

The Geometry of Synchronization Problems and Learning Group Actions

Tingran Gao *

Jacek Brodzki †

Sayan Mukherjee ‡

Abstract

We develop a geometric framework that characterizes the *synchronization problem* — the problem of registering or aligning a collection of objects in a consistent manner. The theory we formulate characterizes the cohomological nature of synchronization based on the classical theory of fibre bundles. We first establish the correspondence between synchronization problems in a topological group G over a connected graph Γ and the moduli space of flat principal G -bundles over Γ , and develop a discrete analogy of the renowned theorem of classifying flat principal bundles with fix base and structural group using the representation variety. In particular, we show that prescribing an edge potential on a graph is equivalent to specifying an equivalence class of flat principal bundles, of which the triviality of holonomy dictates the synchronizability of the edge potential. We then develop a twisted cohomology theory for associated vector bundles of the flat principal bundle arising from an edge potential, which is a discrete version of the twisted cohomology in differential geometry. This theory realizes the obstruction to synchronizability as a cohomology group of the twisted de Rham cochain complex. We then extend the twisted cohomology theory to a twisted Hodge theory — a fibre bundle analog of the discrete Hodge theory on graphs. The lowest-degree Hodge Laplacian of this twisted Hodge theory recovers a geometric realization of the *graph connection Laplacian* (GCL), a group-valued graph operator studied extensively in synchronization problems. Motivated by our geometric framework, we study the problem of *learning group actions* — partitioning a collection of objects based on the local synchronizability of pairwise correspondence relations. A dual interpretation is to learn finitely generated subgroups of an ambient transformation group from noisy observed group elements. A synchronization-based algorithm is also provided, and we demonstrate its efficacy using simulations and real data.

1 Introduction

Over the past century, concepts from differential geometry have had a strong impact on probability theory, statistical inference, and machine learning [Che66, Cra46, Fis22, Mal36, Rao45]. Two central geometric concepts used in these fields have been differential operators (e.g. the Laplace-Beltrami operator [BN04]) and Riemannian metrics (e.g. Fisher information [Fis22]). In particular, the research program of manifold learning studies dimension reduction through the lens of differential-geometric quantities and invariants, and designs data compression algorithms that preserve intrinsic geometric information such as geodesic distances [TdSL00], affine connections [RS00], second fundamental forms [DG03], and heat kernels [BN03, CLL+05a, CLL+05b]. The underlying hypothesis of these techniques is that the data lie approximately on a smooth manifold (often embedded in an ambient Euclidean space), a scenario facilitating inference due to smoothly controllable transitions between observed and unseen data. For practical purposes, discrete analogues of the inherently smooth theory of differential geometry have also been explored in fields ranging from geometry processing [BSSZ08, CdGDS13], finite element methods [AFW06], to spectral graph theory [Chu97] and diffusion geometry [CL06, SW12].

Beyond the manifold assumption, geometric objects can be handled with “softer” tools such as topology: topological data analysis techniques [CZCG05, EH10] have been developed to study datasets based on their persistent homology. For smooth manifolds, it is well known that the singular cohomology and de Rham cohomology are isomorphic, indicating that some topological information can be read off from the differential structure of geometric

*Department of Mathematics, Duke University, Durham, NC 27708, USA (trgao10@math.duke.edu). This author was supported by Simons Math+X Investigators Award #400837.

†Mathematical Sciences, University of Southampton, Southampton SO17 1BJ, UK (j.brodzki@soton.ac.uk). This author was partially supported by the EPSRC grants EP/I016945/1 and EP/N014189/1.

‡Departments of Statistical Science, Mathematics, and Computer Science, Duke University, Durham, NC 27708, USA (sayan@stat.duke.edu). This author was supported by NSF DMS 16-13261, NSF IIS 1546331, NSF DMS-1418261, NSF IIS-1320357, and NSF DMS-1045153.

objects. Carrying the de Rham theory beyond the manifold setting has attracted the interest of geometers and physicists: synthetic differential geometry [Koc82, Koc06] defines group-valued differential forms on “formal manifolds” (generalized notion of smooth spaces for which infinitesimal neighborhoods are specified axiomatically), based on which an analog of the classical de Rham theory can be established [FL90]; noncommutative differential geometry [Con85, Con00, Mad99] builds upon the observation that much of differential geometry can be formulated in terms of the algebra of smooth functions defined on smooth manifolds, and replaces this algebra with noncommutative ones — differential forms can then be extended to “noncommutative spaces” along with homology and cohomology of much more general objects. Discrete analogs of the Hodge Laplacian, a second order differential operator closely related to de Rham theory, have been proposed for simplicial complexes and graphs [JLYY11, Lim15, PR16, PRT15, SKM14]; its non-commutative counterpart for 1-forms on graphs have recently been explored in [Maj13].

Bridging recent developments applying differential geometry and topology in probability and statistical sciences, the problem of *synchronization* [BSS13, WS13] arise in a variety of fields in computer science (e.g. computer vision [BCSZ14] and geometry processing [KKBL15]), signal processing (e.g. sensor network localization [CLS12]), combinatorial optimization (e.g. non-commutative Grothendieck inequality [BKS16]), and natural sciences (e.g. cryo-electron microscopy [BCS15, SS12, SZSH11] and geometric morphometrics [Gao15]). The data given in a synchronization problem include a connected graph that encodes similarity relations within a collection of objects, and pairwise correspondences — often realized as elements of a transformation group G — characterizing the nature of the similarity between a pair of objects linked directly by an edge in the relation graph. The general goal of the problem is to adjust the pairwise correspondences, which often suffer from noisy or incomplete measurements, to obtain a globally consistent characterization of the pairwise relations for the entire dataset, in the sense that unveiling the transformation between a pair of objects far-apart in the relation graph can be done by composing transformations along consecutive edges on a path connecting the two objects, and the resulting composed transformation is independent of the choice of the path. (A precise definition of a synchronization problem will be provided below; see Section 1.1.) This paper stems from our attempt to gain a deeper understanding of the geometry underlying synchronization problems, and its relation to the problem of learning group actions. Whereas previous works [SW12, SW13] in this direction build upon manifold assumptions, the point of view we adopt here is synthetic and noncommutative: we will see that inference is possible due to rigidity rather than smoothness.

The remainder of this section gives a formal definition of synchronization problems, as well as a geometric interpretation in the language of fibre bundles. The fibre bundle interpretation is elementary but has not been presented in the literature of synchronization problems, to our knowledge. We then state the main results, discuss related works, and describe the organization of the paper.

1.1 A Fibre Bundle Interpretation of Synchronization Problems

We begin with a standard formulation of the synchronization problem originated in a series of works by A. Singer and collaborators [Sin11, SS12, BSS13, WS13, BSAB14, BCS15]. Let $\Gamma = (V, E, w)$ be an undirected weighted graph with vertex set V , edge set E , and weights w_{ij} for each $(i, j) \in E$. Assume G is a topological group acting on a normed vector space F . Given a map $\rho : E \rightarrow G$ from the edges of Γ to the group G satisfying $\rho_{ij} = \rho_{ji}^{-1}$, the objective of a *F-synchronization problem over Γ with respect to ρ* is to find a map $f : V \rightarrow F$ satisfying the constraints

$$f_i = \rho_{ij} f_j \quad \forall (i, j) \in E. \quad (1)$$

If no such map f exists, the synchronization problem consists of finding a map f from V to F that satisfies the constraints as much as possible, in the sense of minimizing the *frustration*

$$\eta(f) = \frac{1}{2} \frac{\sum_{i,j \in V} w_{ij} \|f_i - \rho_{ij} f_j\|_F^2}{\sum_{i \in V} d_i \|f_i\|_F^2}, \quad (2)$$

where $\|\cdot\|_F$ is a norm defined on F , and $d_i = \sum_{j:(i,j) \in E} w_{ij}$ is the weighted degree at vertex i . In the terminology of [BSS13], ρ is an *edge potential* and f is a *vertex potential*; a vertex potential is said to *satisfy* a given edge potential if all equalities in (1) hold. Varying the choice of group G and field F results in different realizations of the synchronization problem [SW11, SZSH11, TSR11, SS12, WS13, BSS13, BSAB14], as will be elaborated in Section 1.3.

Table 1: Notations used throughout this paper

Γ	graph
V	vertex set of Γ
E	edge set of Γ
n or $ V $	number of vertices of the graph Γ
m or $ E $	number of edges of the graph Γ
w_{ij}	weight on edge $(i, j) \in E$
d_i	weighted degree on vertex $i \in V$, defined as $d_i = \sum_{j:(i,j) \in E} w_{ij}$
G	topological group
e	identity element of G
G_δ	group G equipped with discrete topology
\mathbb{K}	scalar field \mathbb{R} or \mathbb{C}
F	vector space on \mathbb{K} that is a representation space of G
d or $\dim F$	dimension of the vector space F
$\langle \cdot, \cdot \rangle_F$	inner product on F
$\mathfrak{U} = \{U_i \mid 1 \leq i \leq V \}$	open cover of Γ in which U_i is the star of vertex $i \in V$
$C^0(\Gamma; G)$	G -valued 0-cochain on Γ , or the set of all vertex potentials on Γ
$C^1(\Gamma; G)$	G -valued 1-cochain on Γ , or the set of all edge potentials on Γ
$C^0(\Gamma; F)$	F -valued 0-cochain on Γ
\mathcal{B}_ρ	synchronization principal bundle (a flat principal G -bundle on Γ) associated with $\rho \in C^1(\Gamma; G)$
$\mathcal{B}_\rho[F]$	flat associated F -bundle of \mathcal{B}_ρ
hol_ρ	holonomy homomorphism on \mathcal{B}_ρ , from $\pi_1(\Gamma)$ to G
$\text{Hol}_\rho(\Gamma)$	holonomy of the synchronization principal bundle \mathcal{B}_ρ
$\Omega_i^0(\Gamma; \mathcal{B}_\rho[F])$	constant twisted local 0-forms of $\mathcal{B}_\rho[F]$ on U_i , i.e. constant local sections of $\mathcal{B}_\rho[F]$ on U_i
$\Omega^0(\Gamma; \mathcal{B}_\rho[F])$	locally constant twisted global 0-forms of $\mathcal{B}_\rho[F]$, i.e. locally constant global sections of $\mathcal{B}_\rho[F]$
$\Omega_i^1(\Gamma; \mathcal{B}_\rho[F])$	constant twisted local 1-forms of $\mathcal{B}_\rho[F]$ on U_i
$\Omega^1(\Gamma; \mathcal{B}_\rho[F])$	locally constant twisted global 1-forms of $\mathcal{B}_\rho[F]$
$[f]$	vector in \mathbb{K}^{nd} representing $f \in C^0(\Gamma; F)$
d_ρ	ρ -twisted differential where $\rho \in C^1(\Gamma; G)$, from $C^0(\Gamma; F)$ to $\Omega^1(\Gamma; \mathcal{B}_\rho[F])$
δ_ρ	ρ -twisted codifferential where $\rho \in C^1(\Gamma; G)$, from $\Omega^1(\Gamma; \mathcal{B}_\rho[F])$ to $C^0(\Gamma; F)$
$\Delta_\rho^{(0)}$	ρ -twisted Hodge Laplacian of degree 0
$\Delta_\rho^{(1)}$	ρ -twisted Hodge Laplacian of degree 1
$H_\rho^0(\Gamma; \mathcal{B}_\rho[F])$	The 0th twisted cohomology group for $\mathcal{B}_\rho[F]$, where $\rho \in C^1(\Gamma; G)$
$\nu(S)$	frustration of the subgraph of Γ spanned by the vertex subset $S \subset V$

Since we will frequently refer to the set of all edge and vertex potentials on a graph, let us introduce the following notations to ease our exposition: let $C^0(\Gamma; G)$, $C^1(\Gamma; G)$ denote respectively the set of all G -valued vertex and edge potentials on Γ , i.e.

$$C^0(\Gamma; G) := \{f : V \rightarrow G\}, \quad C^1(\Gamma; G) := \{\rho : E \rightarrow G \mid \rho_{ij} = \rho_{ji}^{-1}, \forall (i, j) \in E\}. \quad (3)$$

For cohomological reasons that will become clear in Section 2, we will also call $C^0(\Gamma; G)$ and $C^1(\Gamma; G)$ the G -valued 0- and 1-cochains on Γ , respectively. Similarly, let

$$C^0(\Gamma; F) := \{f : V \rightarrow F\} \quad (4)$$

denote the set of all F -valued vertex potentials on Γ . Throughout this paper, a G -valued edge potential $\rho \in \Omega^0(\Gamma; G)$ is said to be *synchronizable* if there exists a G -valued vertex potential $f \in \Omega^0(\Gamma; G)$ satisfying $f_i = \rho_{ij} f_j$, $\forall (i, j) \in E$, i.e. (1) is satisfied with $F = G$. Generally, an F -valued vertex potential satisfying (1) will be referred to as a *solution to the F -synchronizable problem over Γ with respect to ρ* , or simply *F -valued synchronization solution*. Clearly, ρ is synchronizable if and only if a G -valued synchronization solution exists.

When $F = G$, i.e. when we consider the action of G on itself, a synchronizable edge potential can be realized geometrically as a *flat*¹ principal bundle that is isomorphic to a product space in its entirety, i.e. a *trivial*² flat principal bundle, as will be explained in Proposition 1.1 and Proposition 1.2 below; this observation forms the backbone of the entire geometric framework we develop in this paper. When the fibre bundle is differentiable, this notion of flatness is equivalent to the existence of a flat connection on the bundle, which is essentially a special case of the Riemann-Hilbert correspondence [Esn88]. The main results of this paper build upon extending further and deeper the analogy between the geometry of synchronization problems and fibre bundles.

Propositions 1.1 and 1.2 characterize the basic building block for the geometric formulation of synchronization problems. We will develop the principal bundle in the generality of topological spaces that includes smooth structures as particular cases. Following Steenrod [Ste51], a fibre bundle is a quintuple $\mathcal{E} = (E, M, F, \pi, G)$ where E , M , F are topological spaces, referred to as the *total space*, *base space*, and *fibre space*, respectively; $\pi : E \rightarrow M$ is a continuous surjective map, called the *bundle projection*, and M adopts an open cover $\{U_i\}$ with homeomorphisms $\phi_i : U_i \times F \rightarrow \pi^{-1}(U_i)$ between each $\pi^{-1}(U_i) \subset E$ and the product space $U_i \times F$, such that $\pi|_{\pi^{-1}(U_i)}$ is the composition of ϕ_i with $\text{proj}_1 : U_i \times F \rightarrow U_i$, the canonical projection onto the first factor of the product space. In other words, the following diagram is commutative:

$$\begin{array}{ccc} \pi^{-1}(U_i) & \xrightarrow{\phi_i} & U_i \times F \\ & \searrow \pi & \swarrow \text{Proj}_1 \\ & U_i & \end{array}$$

The open cover $\{U_i\}$ and the homeomorphisms $\{\phi_i\}$ together provides a system of *local trivializations* for the fibre bundle \mathcal{E} . Moreover, G is a topological transformation group on F encoding the compatibility of “change-of-coordinates” on M , with respect to the provided local trivializations, in the following sense: at every $x \in U_i \cap U_j \neq \emptyset$, the restriction of the composed map $\phi_i^{-1} \circ \phi_j : U_j \times F \rightarrow U_i \times F$ on $\{x\} \times F$, which necessarily gives rise to a homeomorphism from $\{x\} \times F$ to itself by definition, is canonically identified with a group element $g_{ij}(x) \in G$, and the map $g_{ij} : U_i \cap U_j \rightarrow G$ is continuous. The topological group G is called the *structure group* of the fibre bundle \mathcal{E} . The notation F_x is often used to denote $\pi^{-1}(x)$ for $x \in M$, and referred to as the *fibre over $x \in M$* . It is straightforward to check from these definitions that

$$g_{ii}(x) = e \quad \forall x \in U_i \quad (5)$$

$$g_{ij}(x) = g_{ji}^{-1}(x) \quad \forall x \in U_i \cap U_j \quad (6)$$

$$g_{ij}(x) g_{jk}(x) = g_{ik}(x) \quad \forall x \in U_i \cap U_j \cap U_k \quad (7)$$

¹Recall (see, e.g. [Tu83, §2]) that a fibre bundle $\pi : \mathcal{B} \rightarrow X$, with total space \mathcal{B} and base space X , is said to be flat if it admits a system of local trivializations with locally constant bundle coordinate transformations.

²Note that a flat bundle is not necessarily trivial (i.e. isomorphic to a product space) — the fundamental group of the base space plays a central role in this development (see e.g. [Mor01, Chapter 2]).

where e is the identity element of the structural group G . The family of continuous maps $\{g_{ij} : U_i \cap U_j \rightarrow G\}$ is called a system of *coordinate transformations for the fibre bundle* \mathcal{E} . Interestingly, essentially all information for determining the fibre bundle \mathcal{E} is encoded in the coordinate transformations, as the following theorem indicates:

Theorem 1.1 (Steenrod [Ste51] §3.2). *If G is a topological transformation group of F , U_j is an open cover of M , $\{g_{ij}\}$ is a family of continuous maps from each non-empty intersection $U_i \cap U_j$ to G satisfying (5), (6), (7), then there exists a fibre bundle \mathcal{E} with base space M , fibre F , structural group G , and coordinate transformations $\{g_{ij}\}$. Any two such fibre bundles are equivalent to each other.*

The precise definition for two fibre bundles with the same base space, fibre space, and structural group to be equivalent can be found in [Ste51, §2.4], but we will also cover it in Section 2.1. Notice that the conditions (5), (6), (7) are reminiscent of the characterization for the synchronizability (1) of a G -valued edge potential on a connected graph: if ρ satisfies (1) for a map $f : V \rightarrow G$, then $\rho_{ij} = f_i f_j^{-1}$ on each edge $(i, j) \in E$, which certainly satisfies

$$\rho_{ii} = e \quad \forall i \in V, \quad \rho_{ij} = \rho_{ji}^{-1} \quad \forall (i, j) \in E, \quad \rho_{ij} \rho_{jk} = \rho_{ik} \quad \forall (i, j), (j, k), (i, k) \in E. \quad (8)$$

As the following Proposition 1.1 establishes, viewing the graph Γ as a topological space, an appropriate open cover of Γ can be found such that any synchronizable edge potential can be realized as coordinate transformations of a fibre bundle with base space Γ and the topological group G serving both as the fibre space and the structure group. A fibre bundle with its structural group as fibre type is called a *principal bundle*. Moreover, any such principal bundle must also be flat, as the bundle coordinate transformations take constant values on every non-empty intersection of sets in the open cover. The following simple concepts from combinatorial graph theory and algebraic topology (see, e.g. [BC08, BT82]) will be needed for the statement and proof of Proposition 1.1:

- 1) A *clique complex* of a graph $\Gamma = (V, E)$ is the simplicial complex with all complete subgraphs of Γ as its faces;
- 2) The *star* of a vertex v in a simplicial complex \mathcal{K} is the union of all closed simplices in \mathcal{K} containing v as a vertex;
- 3) The *support* of a simplicial complex \mathcal{K} is the underlying topological space of \mathcal{K} .

Proposition 1.1. *Let G be a topological group, $\Gamma = (V, E)$ a connected undirected graph, and $\rho : E \rightarrow G$ a map satisfying $\rho_{ij} = \rho_{ji}^{-1}$ for all $(i, j) \in E$. Denote X for the 2-skeleton of the clique complex of the graph Γ , \mathcal{X} the support of X , and $\mathfrak{U} = \{U_i \mid 1 \leq i \leq |V|\}$ for an open cover of \mathcal{X} in which U_i is the interior of the star of vertex i . Then ρ is synchronizable over G if and only if there exists a flat trivial principal fibre bundle $\pi : \mathcal{P}_\rho \rightarrow \mathcal{X}$ with structure group G and a system of local trivializations defined on the open sets in \mathfrak{U} with constant bundle transition functions ρ_{ij} on non-empty $U_i \cap U_j$.*

A proof of Proposition 1.1 can be found in Appendix A. The key idea is to view Γ as the 1-skeleton of its associated clique complex, and use the open cover consisting of star neighborhoods of each vertex. A similar construction of “Cryo-EM complex” has been used in [YL16] to classify data input to Cryo-EM problems, an important application of synchronization techniques.

However, it is important to notice that the converse to Proposition (1.1) is not true in general; more precisely, an edge potential satisfying (8), which necessarily specifies a flat principal bundle over Γ , need not be synchronizable. For a simple example, consider a square graph Γ consisting of a four vertices 1, 2, 3, 4 and four edges (1, 2), (2, 3), (3, 4), (4, 1), forming a closed simple loop but without any triangles enclosed by three edges. An edge potential satisfying $\rho_{ij} = \rho_{ji}^{-1}$ on all edges clearly satisfies all equalities in (8) since no consistency needs to be checked on edge triplets, but it is easy to find ρ violating the equality $\rho_{12}\rho_{23}\rho_{34}\rho_{41} = e$ which must be obeyed by any synchronizable edge potential, provided that the group G is not trivial. The lesson is that the compatibility conditions (8) are of a local nature, in the sense that the *cycle-consistency* (borrowing a term from geometry processing of shape collections [NBCW⁺11a, HG13] that describes a compatibility constraint analogous to the last equality in (8)) is imposed only on triangles composed of edge triplets; in contrast, synchronizability requires a stronger notion of “global” cycle-consistency for the operation of composing group elements along loops of arbitrary length and topology on the graph. In a certain sense, fibre bundles are the geometric models realizing edge potentials that are “locally synchronizable.”

Proposition 1.1 is our first attempt at understanding the geometric mechanism of synchronization problems. The assumption of the synchronizability of ρ significantly restricts the range of applicability of this geometric analogy: in most scenarios of interest, the synchronizability of an edge potential is the goal rather than the starting point for a synchronization problem. Fortunately, it is possible to extend the fibre bundle analogy beyond the synchronizability

assumption in Proposition 1.1, by restricting the model base space from the 2-skeleton of the clique complex of the graph to the 1-skeleton, and adjust the open cover \mathfrak{U} accordingly: if we define an open cover \mathfrak{U} on the graph Γ (which as a topological space is canonically identified with the 1-skeleton of its clique complex) in which each open set U_i covers only vertex i and the interior of all edges adjacent to it, then $U_i \cap U_j \neq \emptyset$ if and only if $(i, j) \in E$, and any triple intersection of open sets in \mathfrak{U} is empty. In consequence, any system of bundle coordinate transformations defined on \mathfrak{U} by a G -valued edge potential ρ automatically satisfies (5), (6), (7), and specifies a flat principal G -bundle over Γ , denoted as \mathcal{B}_ρ , regardless of synchronizability. This is also consistent with the definition of vector bundles on graphs in [Ken11]. Clearly, when ρ is synchronizable, \mathcal{B}_ρ is the restriction of the principal G -bundle \mathcal{P}_ρ in Proposition 1.1 to the 1-skeleton of the base space Γ , therefore trivial as well. Conversely, if \mathcal{B}_ρ is trivial, by [Ste51, §2.10 or §4.3], there exists a map $f : \Gamma \rightarrow G$ assigning a constant value f_i for all points $x \in U_i$ such that $\rho_{ij} = f_i f_j^{-1}$ for all $U_i \cap U_j \neq \emptyset$, which gives rise to a map $f : V \rightarrow G$ by restriction to the vertex set V of Γ ; this verifies all constraints in (1) and establishes the synchronizability of the edge potential ρ . Consequently, the triviality of \mathcal{B}_ρ and \mathcal{P}_ρ implies each other, both are equivalent to the synchronizability of ρ . We summarize these observations in Proposition 1.2 and formally define the *synchronization principal bundle* \mathcal{B}_ρ , which will be of central importance for the geometric framework we develop in the rest of this paper.

Proposition 1.2. *Let G be a topological group, $\Gamma = (V, E)$ a connected undirected graph, and $\rho : E \rightarrow G$ a map satisfying $\rho_{ij} = \rho_{ji}^{-1}$ for all $(i, j) \in E$. Write $\mathfrak{U} = \{U_i \mid 1 \leq i \leq |V|\}$ for an open cover of Γ in which U_i is the union of the single vertex set $\{i\}$ with the interior of all edges adjacent to the vertex i . Then ρ defines a flat principal G -bundle \mathcal{B}_ρ over Γ with a system of local trivializations defined on the open sets in \mathfrak{U} with constant bundle transition functions ρ_{ij} on non-empty $U_i \cap U_j$. Furthermore, ρ is synchronizable if and only if \mathcal{B}_ρ is trivial.*

Definition 1.1 (Synchronization Principal Bundle). The fibre bundle \mathcal{B}_ρ associated with the connected graph Γ and edge potential ρ as characterized in Proposition 1.2 is called a *synchronization principal bundle of edge potential ρ over Γ* , or a *synchronization principal bundle* for short.

In practice, it is often more convenient to work with \mathcal{B}_ρ rather than \mathcal{P}_ρ , not only since non-synchronizable edge potentials are much more prevalent, but also because noisy or incomplete measurements almost always cause the observed group elements $\rho_{ij} \in G$ to be non-synchronizable. Solving for a G -valued synchronization solution can thus be viewed as an approach to “denoising” or “filtering” those observed transformations ρ_{ij} , as was already implicit in many applications [BSS13, BKS16]. In the sense of Proposition 1.2, these problems can be interpreted as inference on the structure of flat principal bundles.

Most synchronization problems in practice [SW11, SZSH11, SS12, TSR11] considers vertex potentials valued in G , the same topological group in which the prescribed edge potential takes value, pertaining to the principal bundle picture (i.e. $F = G$) discussed in Proposition 1.1 and Proposition 1.2. Our fibre bundle interpretation naturally includes more general synchronization problems in which the vertex potentials takes values in $F \neq G$ as well, by relating the synchronizability of an edge potential to the existence of global sections on an *associated F -bundle* $\mathcal{P}_\rho \times_\eta F$ or $\mathcal{B}_\rho \times_\eta F$, where $\eta : G \times F \rightarrow F$ denotes the action of G on F . Essentially, an associated F -bundle $\mathcal{P}_\rho \times_\eta F$ (or $\mathcal{B}_\rho \times_\eta F$) is constructed using the same procedure as the principal bundle \mathcal{P}_ρ (or \mathcal{B}_ρ), but with fibre F instead of G . The associated bundles $\mathcal{P}_\rho \times_\eta F$, $\mathcal{B}_\rho \times_\eta F$ are thus also flat (but not necessarily trivial), since their bundle coordinate transformations are equivalent to those of their principal bundle, up to the representation induced by the action η .

A major difference between working with an associated bundle and the principal bundle is that the cocycle condition $\rho_{kj}\rho_{ji} = \rho_{ki}$ may still not be satisfied in the presence of a vertex potential $f : V \rightarrow F$ satisfying (1), as elements in F can not be “inverted” in general; another important difference is the relation between triviality and global sections: whereas a principal bundle \mathcal{P}_ρ or \mathcal{B}_ρ is trivial if and only if *one* global section exists, which amounts to finding *one* solution to the synchronization problem over Γ with respect to ρ , an associated bundle may admit one or more global sections yet still be non-trivial. For instance, a vector bundle always admits the zero global section, regardless of its triviality. It turns out that establishing the synchronizability of an edge potential through the triviality of an associated bundle requires finding “sufficiently many” global sections of an associated bundle, or in terms of synchronization problems, “sufficiently many” solutions satisfying (1). This is also reflected in the twisted Hodge theory we develop in Section 2.2. Even though finding enough global sections seems to be more work, in practice it could be much easier to find a set of global sections on the associated bundle than to find even only one global section on the principal bundle, as the action of G on the space F introduces additional information from both geometric and practical points of view. Since the flat F -bundle associated with \mathcal{B}_ρ will be essential for Section 2.2, we introduce the following definition:

Definition 1.2 (Synchronization Associated Bundle). The flat F -bundle on Γ associated with the flat principal bundle

\mathcal{B}_ρ by the action of G on F is called a *synchronization associated F -bundle of edge potential ρ over Γ* , or *synchronization associated bundle* for short, denoted as $\mathcal{B}_\rho[F]$.

We close this preliminary section drawing analogy between synchronization problems and fibre bundles by clarifying the relation among, and the geometric implications of, some variants of the optimization formulation of synchronization problems. Given a graph $\Gamma = (V, E)$ and a G -valued edge potential ρ , a direct translation of the goal of finding an F -valued vertex potential f satisfying (1) as much as possible is to solve

$$\min_{f: V \rightarrow F} \sum_{(i,j) \in E} \text{Cost}_F(\rho_{ij} f_j, f_i), \quad (9)$$

where $\text{Cost}_F : F \times F \rightarrow [0, \infty)$ is a cost function on F (e.g., derived from a distance or a norm). When we seek multiple solutions to an F -synchronization problem over Γ with prescribed edge potential ρ , it is natural to impose the additional constraints that the solutions should be sufficiently different from each other; in the presence of a Hilbert space structure on F , it is convenient to impose orthogonality constraints between pairs of solutions to obtain linearly independence. With additional normalization constraints to fix the issue of identifiability, this exactly translates into the spectral relaxation algorithm in [BSS13]. If the prescribed edge potential is synchronizable, its synchronizability will be confirmed once a sufficient number of synchronization solutions can be collected, where the actual number depends on the property of the group G as well as its action on F ; if not, a synchronizable edge potential can be constructed from sufficiently many F -valued “approximate solutions” that minimize the objective function in (9) as much as possible. The case $G = F$ corresponds to the optimization problem

$$\min_{f: V \rightarrow G} \sum_{(i,j) \in E} \text{Cost}_G(\rho_{ij} f_j, f_i). \quad (10)$$

which, in the case the Cost_G is G -invariant, is equivalent to

$$\min_{f: V \rightarrow G} \sum_{(i,j) \in E} \text{Cost}_G(\rho_{ij}, f_i f_j^{-1}). \quad (11)$$

If ρ is synchronizable, a minimizer of (10) (or (9)) attaining zero objective value can be geometrically realized as a global section of the synchronization principal bundle \mathcal{B}_ρ (or \mathcal{P}_ρ); such a minimizer implies the triviality of the principal bundle \mathcal{B}_ρ (or \mathcal{P}_ρ), but not necessarily so in general for the associated bundle $\mathcal{B}_\rho \times_\eta F$ (or $\mathcal{P}_\rho \times_\eta F$). If ρ is not synchronizable, the minimum values of (10), (11), and (9) are all greater than zero, and minimizer of (10) or (11) can be viewed as a “denoised” or “filtered” version of a trivial flat principal bundle underlying the dataset.

1.2 Main Contributions

In this section we give a brief overview of our main contribution. We will motivate the two ingredients of the geometric framework developed in Section 2, namely, holonomy representation and Hodge theory, by demonstrating preliminary versions of our formulation that lead to weaker conclusions or incomplete geometric pictures, then sketch the full approaches adopted in Section 2. Finally, we draw the link between the geometric framework and the proposal of the *learning group actions* (LGA) problem.

1.2.1 Holonomy of Synchronization Principal Bundles

Consider \mathcal{B}_ρ , the synchronization principal bundle arising from a G -synchronization problem over a connected graph $\Gamma = (V, E)$ with respect to $\rho \in C^1(\Gamma; G)$. Fix an arbitrary vertex $v \in V$, and denote the set of all v -based loops in Γ (loops with v as both the starting and ending vertex) as Ω_v ; Ω_v carries a natural group structure by the composition of v -based loops. Now the procedure of taking the product of the values of ρ along the consecutive edges in the loop specifies a group homomorphism from Ω_v into G . Denote the image of this group homomorphism by H_v , which is necessarily a finitely generated subgroup of G since Γ is a finite graph. The simple but important observation here is that H_v is the trivial subgroup of G if and only if ρ is synchronizable. The group H_v is the analogy of the *holonomy group based at v* in differential geometry, if we view ρ_{ij} on edge (i, j) as the *parallel-transport* between fibres of \mathcal{B}_ρ at $i, j \in V$.

Section 2.1 is devoted to a deeper and more systematic treatment of the group homomorphism from loops in Γ to the structure group G . We will define $\text{Hol}_\rho(\Gamma)$, the holonomy of the synchronization principal bundle \mathcal{B}_ρ (independent of the choice of the base vertex v), as well as an equivalence relation on $C^1(\Gamma; G)$ induced by a right action of $C^0(\Gamma; G)$ (which is treated implicitly when solving synchronization problems in practice), and establish a correspondence between $\text{Hol}_\rho(\Gamma)$ and the equivalence class in $C^1(\Gamma; G)/C^0(\Gamma; G)$ to which $\rho \in C^1(\Gamma; G)$ belongs. In particular, trivial holonomy $\text{Hol}_\rho(\Gamma)$ corresponds to the orbit in $C^1(\Gamma; G)/C^0(\Gamma; G)$ consisting precisely of all synchronizable edge potentials. This correspondence will be formulated in Theorem 2.1 as between $C^1(\Gamma; G)/C^0(\Gamma; G)$ and $\text{Hom}(\pi_1(\Gamma), G)/G$, the *representation variety* of the fundamental group of Γ into G .

1.2.2 Twisted De Rham Cohomology of Associated Vector Bundles

The *graph connection Laplacian* (GCL) for an F -synchronization problem over graph Γ with respect to $\rho \in \Omega^1(\Gamma; G)$ is a linear operator on $\Omega^0(\Gamma; F)$ defined as

$$(L_1 f)_i := \frac{1}{d_i} \sum_{j: (i,j) \in E} w_{ij} (f_i - \rho_{ij} f_j), \quad \forall i \in V, \quad \forall f \in C^0(\Gamma; F).$$

If F is a vector space and G has a matrix representation on F , GCL can be written as a block matrix in which the (i, j) -th block is the matrix representation of ρ_{ij} , if $(i, j) \in E$. GCL essentially carries all information of a synchronization problem and is of central importance to the spectral and SDP relaxation algorithms for synchronization. Our motivation for Section 2.2 was to provide a cohomological interpretation for GCL, in the hope of realizing it geometrically as a Hodge Laplacian in a cochain complex, inspired by a similar geometric realization of the *graph Laplacian* (in the context of algebraic and spectral graph theory) as a Hodge Laplacian on degree-zero forms in discrete Hodge theory (see Appendix B). Note that GCL reduces to the graph Laplacian if the group G is a scalar field.

In the literature of differential geometry, *twisted differential forms* on a flat vector bundle \mathcal{E} can be intuitively thought as bundle-valued differential forms on the base manifold. The twisted Hodge theory we develop in Section 2.2 defines two discrete differential operators that are formal adjoints of each other between *constant twisted local 0-forms* and *constant twisted local 1-forms* on the synchronization associated F -bundle $\mathcal{B}_\rho[F]$, namely the ρ -*twisted differential* d_ρ and the ρ -*twisted codifferential* δ_ρ , such that GCL can be written as the composition $\delta_\rho d_\rho$. Provisionally, by identifying each $f \in C^0(\Gamma; F)$ naturally with a collection of constant twisted local 0-forms — one for each open set $U_i \in \mathcal{U}$ — a coarse approximation of our construction can be written as

$$\begin{aligned} (d_\rho f)_{ij} &\sim f_i - \rho_{ij} f_j, \quad \forall f \in C^0(\Gamma; F), \\ (\delta_\rho \omega)_i &\sim \frac{1}{d_i} \sum_{j: (i,j) \in E} w_{ij} \omega_{ij}, \quad \forall \omega \in C^1(\Gamma; F) := \left\{ \omega : E \rightarrow F \mid \omega_{ij} = -\omega_{ji} \quad \forall (i, j) \in E \right\}, \end{aligned}$$

from which it can be easily checked that $L_1 = \delta_\rho d_\rho$ on $C^0(\Gamma; F)$. The main conceptual difficulty with this natural formulation is that “ $d_\rho f$ ” defined as such does not possess the skew-symmetry desired for 1-forms, since in general

$$f_j - \rho_{ji} f_i = -\rho_{ji} (f_i - \rho_{ij} f_j) \neq -(f_i - \rho_{ij} f_j). \quad (12)$$

The framework we develop in Section 2.2 circumvent this skew-symmetry issue with 1-forms by defining $f_i - \rho_{ij} f_j$ as the representation of $d_\rho f$, a *twisted global 1-form* defined over the entire graph Γ , in the system of local trivializations of $\mathcal{B}_\rho[F]$ over the open cover \mathcal{U} . We then define the ρ -*twisted codifferential* δ_ρ that is the formal adjoint of d_ρ with respect to inner products naturally specified on the space of twisted local 0- and 1-forms, and realize the graph connection Laplacian L_1 as the degree-zero Hodge Laplacian $\delta_\rho d_\rho : C^0(\Gamma; G) \rightarrow C^0(\Gamma; G)$ in the twisted de Rham cochain complex (41). These constructions lead to two different characterizations of the synchronizability of ρ , one in Proposition 2.3 with a twisted de Rham cohomology group, and the other through a Hodge-type decomposition of $C^0(\Gamma; F)$ following Theorem 2.3. This twisted Hodge theory also provides geometric insights for the GCL-based spectral relaxation algorithm and Cheeger-type inequalities in [BSS13], as we will elaborate in Section 2.2.4.

1.2.3 Learning Group Actions via Synchronizability

Fibre bundles are topological spaces that are product spaces locally but not necessarily globally. However, we can still look for maximal open subsets of the base space on which the fibre bundle is trivializable, and seek a decomposition

of the base space into the union of such “maximal trivializable subsets.” This intuition motivated us to consider applying synchronization techniques to partition a graph into connected components, based on the synchronizability of a prescribed edge potential in addition to the connectivity of the graph. In Section 3.1, we define the general problem of *learning group actions* (LGA) for a set X , equipped with an action by group G , as searching for a partition of X into a specified number of subsets and learning a new group action on X that is cycle-consistent within each partition; the cycle-consistency need not be maintained for a cycle of actions across multiple partitions. The LGA problem is then specialized to the setting of synchronization problems (*learning group actions by synchronization*, or LGAS), for which we define a quantity that measures the performance of graph partitions based on the synchronizability of a fixed edge potential on the entire graph Γ , motivated by the classical normalized graph cut algorithm. Finally, we propose in Section 3.2 a heuristic algorithm for LGAS, building upon iteratively applying existing synchronization techniques hierarchically and performing spectral clustering on the edge-wise frustration.

1.3 Broader Context and Related Work

The synchronization problem has been studied for a variety of choices of topological groups G and spaces F . Typically these formulations fall into our principal bundle setting and require an underlying manifold structure. We give a brief summary for the correspondences between choices of G and F in our framework and the practical instances in the synchronization literature: In [BSS13] $G = F = O(d)$ and $G = O(d)$, $F = \mathbb{S}^{d-1}$ are studied; the case $G = F = SO(d)$ is examined in [BSAB14, WS13]; orientation detection or $G = F = O(1)$ is considered in [SW11]; cryo-electron microscopy concerns $G = F = SO(2)$ [SS12, SZSH11]; globally aligning three-dimensional scans is the case where $G = F = SO(3)$, and so is [TSR11].

Our formulation of the synchronization problem considers a broader class of geometric structure than what has been proposed in the literature. Specifically, we do not require a manifold assumption (as the problem is modeled on topological spaces), or a principal bundle structure (as we can work with any associated bundle), or compact and/or commutative structure groups. For comparison, the vector and principal bundle framework developed in [SW12, SW13] relies on manifold assumptions for the base, fibre, and total space, as well as an (extrinsic) isometric embedding into an ambient Euclidean space for locally estimating tangent spaces and parallel-transports; similarly for recent work [Gao16] extending this geometric framework to smooth bundles with general fibre types. Both Vector Diffusion Maps (VDM) [SW12] and Horizontal Diffusion Maps (HDM) [Gao16] can be viewed as attempts at combining the idea of synchronization with diffusion geometry [CL06, CLL+05a, CLL+05b]. The geometry underlying the synchronization problem related to cryo-electron microscopy [SS12, SZSH11] can be described using the language of $SO(2)$ -principal bundles, as recently demonstrated in [YL16], with a Čech cohomology approach through Leray’s Theorem which depends essentially on the commutativity of the structure group $SO(2)$, whereas most synchronization problems of practical interest involve noncommutative structure groups. The Non-Unique Games (NUG) and SDP relaxation framework established in [BCS15, BCSZ14] assumes the compactness of the structure group G , and resorts to a compactification procedure that maps a subset of G to another compact group for synchronization problems over non-compact groups such as the Euclidean group in the motion estimation problem in computer vision [MP07, HTDL13].

The graph twisted Hodge theory we develop in Section 2.2 also has ties to recent developments in discrete Hodge theory [JLYY11, Lim15, MS16, PR16, PRT15, SKM14]. In [Ken11] Laplacians on one- and two-dimensional vector bundles on graphs were used to understand the relation between graphs embedded on surfaces and cycle rooted spanning forests, generalizing the relation between spanning trees and graph Laplacians. Variants of the graph Laplacian have been used to relate ranking problems to synchronization problems in [Cuc16, FS15]; a combinatorial Laplacians based on a discrete Hodge theory on directed graphs has been successfully applied to decompose ranking problems and games into “gradient-like” versus “cyclic” components in [JLYY11, CMOP11], and to visualize directed networks in [FAS16]. Discrete Laplacians on simplicial complexes have been proposed, and spectral properties such as Cheeger inequalities and stationary distributions of random walks have been examined in a series of papers [JLYY11, Lim15, MS16, PR16, PRT15, SKM14]. The cocycle conditions (8) are also characterized in geometry processing and computer vision recently for analysis of shape or image collections [NBCW+11b, HG13, WHG13, HWG14], where they are also known as *cycle-consistency conditions*.

The geometric and topological tools we utilize in this paper, namely those involving the topology and geometry of fibre bundles, are covered in most standard textbooks, e.g. [Ste51, Tau11, BT82]. After Milnor’s seminal work on flat connections on a Riemannian manifold [Mil58], the relation between flat bundles and their holonomy homomorphisms became widely known [KT67, Lue76, Gol82, Vas83, Cor88] and is still attracting interests of modern mathematical physicists (e.g. Higgs bundles and representation of the fundamental group [Hit87, Sim92]). In a com-

pletely topological setup, flat bundles can be characterized as fibrations with a homotopy-invariant lifting property (a topological analogue of parallel-transport in differential geometry); essentially the same correspondence between flat bundles and holonomy homomorphisms is already known to Steenrod [Ste51] and referred to as characteristic classes of flat bundles [Dup76, Lue76, Gol82, Esn88]. In a broader context, the correspondence between flat bundles (integrable connections) and local systems (locally constant sheaves) is a special case of the *Riemann-Hilbert correspondence*, a higher-dimensional generalization of Hilbert’s twenty-first problem [Bol90, AB]. This correspondence fostered important developments in algebraic geometry, including D -modules [Kas79, Kas84, Meb80, Meb84] and Deligne’s work on integrable algebraic connections [Del70]. Understanding the representation varieties of the fundamental groups of Riemann surfaces into Lie groups has been of interest to algebraic geometers, geometric topologists, and representation theorists in the past decades [Lab13, Gol06].

The rest of this paper is organized as follows. Section 2 establishes the geometric framework for synchronization problems, relating the synchronizability of an edge potential ρ to (a) the triviality of the holonomy of the flat principal bundle \mathcal{B}_ρ , in Section 2.1; (b) the dimension of the zero-th degree twisted cohomology group of a ρ -twisted de Rham cochain complex, as well as the dimension of the kernel of the zero-th degree twisted Hodge Laplacian, in Section 2.2. Section 3 defines the problem of learning group actions, and proposes SynCut, a heuristic algorithm based on synchronization and graph spectral techniques. Numerical simulations indicating the effectiveness of SynCut is performed on synthetic datasets in Section 3.3 and on a real dataset of a collection of anatomical surfaces in Section 4. A few problems of potential interest are listed in Section 5 for future exploration.

2 Synchronization as a Cohomology Problem

This section concerns two geometric aspects of the synchronization problem. Section 2.1 links the synchronizability of an edge potential to the triviality of the holonomy group of a flat principal bundle. Section 2.2 establishes a discrete twisted Hodge theory that naturally realizes the graph connection Laplacian as the lowest-order Hodge Laplacian of a twisted de Rham cochain complex. The obstruction to synchronizability of an edge potential turns out to be a cohomology group in the twisted de Rham complex; the degeneracy of this cohomology group is reflected in the spectral information of the twisted Hodge Laplacian, which also provides a geometric interpretation for the relaxation techniques used in solving synchronization problems.

2.1 Holonomy and Synchronizability

The two main results in this section, Corollary 2.2 and Theorem 2.1, relate the synchronizability of an edge potential ρ to the triviality of the holonomy group of the synchronization principal bundle \mathcal{B}_ρ . Our motivation is as follows. Recall from Proposition 1.2 that G -synchronization problems on a fixed graph Γ with different edge potentials are in one-to-one correspondence with flat principal G -bundles over Γ , and the synchronizability of an edge potential translates into the triviality of the bundle; whereas the one-to-one correspondence is stated at the level of local coordinates in Proposition 1.2, the triviality of the principal bundle is a property of the equivalence classes of flat principal G -bundles, which suggests the same level of abstraction for synchronizability. As will be precisely stated later in this subsection, (appropriately defined) equivalence classes of edge potentials form the *moduli space* of flat G -bundles on Γ , and a given edge potential is synchronizable if and only if it belongs to the same equivalence class as the trivial edge potential that assigns each edge of Γ the identity element $e \in G$. The *holonomy group*, or the equivalence classes of *holonomy homomorphisms* from the fundamental group of Γ to the structure group G , is a faithful representation of the moduli space of flat principal G -bundles on Γ ; we thus will be able to detect the synchronizability of an edge potential through the triviality of the associated holonomy group. This argument is reminiscent of classical classification theorems of (1) principal bundles with disconnected structure groups in topology (see e.g. [Ste51, §13.9]); (2) flat connections in differential geometry (see e.g. [Tau11, §13.6]); and (3) holomorphic vector bundles of fixed rank and degree (see e.g. [Wei07, Appendix §2.1]) in complex geometry.

For $f \in C^0(\Gamma; G)$, $\rho \in C^1(\Gamma; G)$, we say that f and ρ are compatible on edge $(i, j) \in E$ if $f_i = \rho_{ij}f_j$, and that f and ρ are compatible on graph Γ if they are compatible on every edge in Γ . Recall from Definition 1.1 that we write \mathcal{B}_ρ for the synchronization principal bundle associated with $\rho \in C^1(\Gamma; G)$, as described in Proposition 1.2. Equivalently, it is often convenient to view \mathcal{B}_ρ as a G_δ -bundle on Γ , where G_δ is the same group as G but equipped with the discrete topology. For $\rho, \tilde{\rho} \in \Gamma$, the G_δ -bundles $\mathcal{B}_\rho, \mathcal{B}_{\tilde{\rho}}$ are equivalent, denoted as $\mathcal{B}_\rho \sim \mathcal{B}_{\tilde{\rho}}$, if a bundle map (see [Ste51, §2.5]) exists between \mathcal{B}_ρ and $\mathcal{B}_{\tilde{\rho}}$ that induces the identity map on the base space Γ . Since $\mathcal{B}_\rho, \mathcal{B}_{\tilde{\rho}}$ have the same base

space, fibre, and structure group, recall from [Ste51, Lemma 2.10] that they are equivalent if and only if there exist continuous functions $\lambda_i : U_i \rightarrow G_\delta$ defined on each $U_i \in \mathfrak{U}$ such that

$$\tilde{\rho}_{ij} = \lambda_i(x)^{-1} \rho_{ij} \lambda_j(x), \quad \forall x \in U_i \cap U_j \neq \emptyset.$$

Since the topology on G_δ is discrete and U_i is connected, λ_i is constant on U_i , and defines a vertex potential by setting $f_i := \lambda_i(v_i)$, where $v_i \in U_i$ is the i th vertex of Γ . This proves the following lemma:

Lemma 2.1. *Two edge potentials $\rho, \tilde{\rho} \in C^1(\Gamma; G)$ define equivalent flat principal G -bundles on Γ if and only if there exists $f \in C^0(\Gamma; G)$ such that*

$$\tilde{\rho}_{ij} = f_i^{-1} \rho_{ij} f_j, \quad \forall (i, j) \in E. \quad (13)$$

In other words, equivalence classes of flat principal G -bundles on Γ determined by edge potentials (through Proposition 1.2) are in one-to-one correspondence with equivalence classes in the orbit space $C^1(\Gamma; G)/C^0(\Gamma; G)$, where the right action of $C^0(\Gamma; G)$ on $C^1(\Gamma; G)$ is defined as

$$[f(\rho)]_{ij} := f_i^{-1} \rho_{ij} f_j, \quad \forall (i, j) \in E. \quad (14)$$

Remark 2.1. *The orbit space $C^1(\Gamma; G)/C^0(\Gamma; G)$ is exactly the first cohomology set $\check{H}^1(\Gamma, \mathfrak{U}, \underline{G})$ for the sheaf of germs of locally constant G -valued functions over Γ with respect to the open cover \mathfrak{U} , where G is possibly nonabelian. It is thus not surprising that the orbit space should identify naturally with isomorphism classes of flat principal G -bundles over Γ (see e.g. [Bry07, Proposition 4.1.2] or [Mic08, §8.1]).*

A *path* in Γ is a collection of consecutive edges in Γ . If all edges in path γ are oriented consistently, we say γ is an oriented path. For any oriented path γ , define γ^{-1} the *reverse* of γ as the path in Γ consisting of the same consecutive edges in γ listed in the opposite order and with all orientations reversed. For an oriented path γ consisting of consecutive edges $\{(i_0, i_1), (i_1, i_2), \dots, (i_{N-1}, i_N)\}$ set

$$\text{hol}_\rho(\gamma) = (\rho_{i_N, i_{N-1}} \rho_{i_{N-1}, i_{N-2}} \cdots \rho_{i_2, i_1} \rho_{i_1, i_0})^{-1} = \rho_{i_0, i_1} \rho_{i_1, i_2} \cdots \rho_{i_{N-2}, i_{N-1}} \rho_{i_{N-1}, i_N} \in G, \quad (15)$$

then hol_ρ maps paths in Γ to elements of group G . For two oriented paths γ, γ' such that the ending vertex of γ coincides with the starting vertex of γ' , define $\gamma \circ \gamma'$ as the oriented path constructed by concatenating γ' with γ . It is then straightforward to verify by definition that

$$\text{hol}_\rho(\gamma^{-1}) = \text{hol}_\rho(\gamma)^{-1}, \quad \text{hol}_\rho(\gamma \circ \gamma') = \text{hol}_\rho(\gamma) \text{hol}_\rho(\gamma'). \quad (16)$$

If an oriented path starts and ends at the same vertex v , we call it an *oriented loop based at vertex v* . Denote Ω_v for all loops based at $v \in V$ in Γ , including the single vertex set $\{v\}$ viewed as the identity loop based at v . Clearly, Ω_v carries a group structure with the loop concatenation and reversion operations. The equalities in (16) ensures $\text{hol}_\rho(\{v\}) = e$ and that $\text{hol}_\rho : \Omega_v \rightarrow G$ is a group homomorphism. Moreover, since graph Γ does not contain any 2-simplices, two oriented loops based at v are homotopic if and only if they differ by a collection of disconnected trees in Γ , in which every tree gets mapped to $e \in G$ under the map hol_ρ ; the map $\text{hol}_\rho : \Omega_v \rightarrow G$ thus descends naturally to a map to G from $\pi_1(\Gamma, v)$, the *fundamental group of Γ based at v* . Unless confusions arise, we shall also denote the descended map as hol_ρ for simplicity of notation. Lemma 2.2 below summarizes these discussions.

Lemma 2.2. *The map $\text{hol}_\rho : \pi_1(\Gamma, v) \rightarrow G$ defined in (15) is a group homomorphism. In particular, the image of this homomorphism is a subgroup of G .*

We will refer to the group homomorphism $\text{hol}_\rho : \pi_1(\Gamma, v) \rightarrow G$ as the *holonomy homomorphism at $v \in \Gamma$* for a G -synchronization problem with prescribed edge potential $\rho \in C^1(\Gamma; G)$. Define the *holonomy group at $v \in \Gamma$* of edge potential ρ as the image

$$\text{Hol}_\rho(v) := \text{hol}_\rho(\pi_1(\Gamma, v)).$$

From a different point of view, Lemma 2.2 assigns an element of $\text{Hom}(\pi_1(\Gamma, v), G)$ to each element of $C^1(\Gamma; G)$, where $\text{Hom}(\pi_1(\Gamma, v), G)$ is the set of group homomorphisms from $\pi_1(\Gamma, v)$ to G .

Lemma 2.3. *If Γ is connected, the holonomy groups $\text{Hol}_\rho(v), \text{Hol}_\rho(w)$ at $v, w \in V$ are conjugate to each other as subgroups of G .*

Proof. Let γ be a path in Γ connecting vertex v to vertex w . The fundamental groups of Γ based at v, w are related by conjugation $\pi_1(\Gamma, v) = \gamma^{-1}\pi_1(\Gamma, w)\gamma$, thus

$$\text{Hol}_\rho(v) = \text{Hol}_\rho(\gamma^{-1}\pi_1(\Gamma, w)\gamma) = \text{Hol}_\rho(\gamma^{-1})\text{Hol}_\rho(\pi_1(\Gamma, w))\text{Hol}_\rho(\gamma) = \text{Hol}_\rho(\gamma)^{-1}\text{Hol}_\rho(w)\text{Hol}_\rho(\gamma).$$

□

Define the *holonomy* of $\rho \in C^1(\Gamma; G)$ on a connected graph Γ as the following conjugacy class (orbit of the action by conjugation) of subgroups of G :

$$\text{Hol}_\rho(\Gamma) := \{g^{-1}\text{Hol}_\rho(v)g \mid \text{for all } g \in G, \text{ and an arbitrarily chosen but fixed vertex } v \in V\}. \quad (17)$$

By Lemma 2.3, the definition of $\text{Hol}_\rho(\Gamma)$ is independent of the choice of a fixed base $v \in V$. We say that the holonomy of $\rho \in C^1(\Gamma; G)$ is *trivial* on a connected graph Γ if $\text{Hol}_\rho(\Gamma)$ contains only the trivial subgroup $\{e\}$ for all $g \in G$. Under the connectivity assumption of Γ , the triviality of the global invariant $\text{Hol}_\rho(\Gamma)$ can be completely determined by its seemingly “local” counterparts; see Lemma 2.4 below. Of course, holonomy is not local in nature, as $\text{Hol}_\rho(v)$ encodes the information of all oriented loops based at vertex v and in principle “touches” the entire space Γ .

Lemma 2.4. *If Γ is connected, the following statements are equivalent:*

- (i) $\text{Hol}_\rho(\Gamma)$ is trivial;
- (ii) $\text{Hol}_\rho(v) = \{e\}$ for some vertex $v \in V$;
- (iii) $\text{Hol}_\rho(v) = \{e\}$ for all vertices $v \in V$.

Similar to the definition of $\text{Hol}_\rho(\Gamma)$ in (17), the fundamental group $\pi_1(\Gamma)$ of a connected graph Γ is also determined by the fundamental group $\pi_1(\Gamma, v_0)$ at any vertex $v_0 \in V$ up to conjugacy classes. Therefore, Lemma 2.2 and Lemma 2.3 together assign to each $\rho \in C^1(\Gamma; G)$ an equivalence class in $\text{Hom}(\pi_1(\Gamma), G)/G$, in which G acts on $\text{Hom}(\pi_1(\Gamma), G)$ by the inner automorphisms of G

$$\phi \mapsto g^{-1}\phi g, \quad \forall \phi \in \text{Hom}(\pi_1(\Gamma), G), g \in G.$$

In other words, Lemma 2.2 and Lemma 2.3 guarantee a well-defined map $\text{Hol} : \Omega^1(\Gamma; G) \rightarrow \text{Hom}(\pi_1(\Gamma), G)/G$. Furthermore, note in equation (15) that Hol is invariant under the right action (14) of $C^0(\Gamma; G)$ on $C^1(\Gamma; G)$, thus Hol naturally descends to a map from $C^1(\Gamma; G)/C^0(\Gamma; G)$ to $\text{Hom}(\pi_1(\Gamma), G)/G$. The space $\text{Hom}(\pi_1(\Gamma), G)/G$ is known as the *representation variety* of the fundamental group of Γ (the free product of a finite number of copies of \mathbb{Z}) into G . To simplify the exposition, we shall use the same notation Hol to denote its quotient map induced by the canonical projection $C^1(\Gamma; G) \rightarrow C^1(\Gamma; G)/C^0(\Gamma; G)$. Theorem 2.1 below establishes the bijectivity of the quotient map.

Theorem 2.1. *If Γ is connected, the map $\text{Hol} : C^1(\Gamma; G)/C^0(\Gamma; G) \rightarrow \text{Hom}(\pi_1(\Gamma), G)/G$ defined as $\text{Hol}([\rho]) = [\text{hol}_\rho]$ is bijective. Moreover, $\text{Hom}(\pi_1(\Gamma), G)/G$ is in one-to-one correspondence with equivalence classes of flat principal G -bundles \mathcal{B}_ρ defined by $\rho \in C^1(\Gamma; G)$.*

Proof. We construct an inverse of Hol from $\text{Hom}(\pi_1(\Gamma), G)/G$ back to $C^1(\Gamma; G)/C^0(\Gamma; G)$. Fix an arbitrary vertex $v_0 \in \Gamma$, and let $\chi : \pi_1(\Gamma, v_0) \rightarrow G$ be a group homomorphism. By the connectivity of Γ , each vertex $v_i \in V$ of Γ is connected to v_0 through an oriented path γ_{0i} ; we orient these paths so they all start at vertex v_0 , and enforce $\gamma_{00} = \{v_0\}$. Assign to each edge $(i, j) \in E$ an element ρ_{ij} of the group G defined by

$$\rho_{ij} := \chi(\gamma_{0i} \circ (i, j) \circ \gamma_{0j}^{-1}). \quad (18)$$

Clearly, $\rho_{ij} = \rho_{ji}^{-1}$ follows from the fact that χ is a group homomorphism; so does $\rho_{ii} = e$ for all vertices $v_i \in V$. Of course, an edge potential ρ defined as in (18) depends on the choice of the oriented paths $\{\gamma_{0i}\}$; this dependence is removed after passing to the orbit space $[\rho] \in C^1(\Gamma; G)/C^0(\Gamma; G)$. In fact, let $\{\tilde{\gamma}_{0i}\}$ be an arbitrary choice of $|V|$ oriented paths connecting v_0 to each vertex of Γ satisfying $\tilde{\gamma}_{00} = \{v_0\}$, then

$$\tilde{\rho}_{ij} = \chi(\tilde{\gamma}_{0i} \circ (i, j) \circ \tilde{\gamma}_{0j}^{-1}) = \chi(\tilde{\gamma}_{0i} \circ \gamma_{0i}^{-1})\chi(\gamma_{0i} \circ (i, j) \circ \gamma_{0j}^{-1})\chi(\gamma_{0j} \circ \tilde{\gamma}_{0j}^{-1}) = \chi(\gamma_{0i} \circ \tilde{\gamma}_{0i}^{-1})\rho_{ij}\chi(\gamma_{0j} \circ \tilde{\gamma}_{0j}^{-1}),$$

i.e., as elements in $C^1(\Gamma; G)$, $\tilde{\rho}$ differs from ρ by an action of the vertex potential $f \in C^0(\Gamma; G)$ defined as

$$f_i := \chi(\gamma_{0i} \circ \tilde{\gamma}_{0i}^{-1}), \quad \forall v_i \in V.$$

Therefore, (18) uniquely specifies an element $[\rho]$ in $C^1(\Gamma; G)/C^0(\Gamma; G)$ for any $\chi \in \text{Hom}(\pi_1(\Gamma), G)$.

It remains to show that $\text{Hol}([\rho])$ differs from χ by an inner automorphism of G . To see this, let ω be an arbitrary oriented loop on Γ based at v_0 consisting of consecutive edges $(v_0, v_{i_1}), (v_{i_1}, v_{i_2}), \dots, (v_{i_N}, v_0)$, where N is some nonnegative integer. Using $\gamma_{00} = \{0\}$ and $\gamma_{0i}^{-1} \circ \gamma_{0i} = \{v_0\}$ for any $v_i \in V$, we have

$$\begin{aligned} \text{hol}_\rho(\omega) &= (\rho_{0,i_N} \rho_{i_N,i_{N-1}} \cdots \rho_{i_2,i_1} \rho_{i_1,0})^{-1} = \rho_{0,i_1} \rho_{i_1,i_2} \cdots \rho_{i_{N-1},i_N} \rho_{i_N,0} \\ &= \chi(\gamma_{00} \circ (v_0, v_{i_1}) \circ \gamma_{0,i_1}^{-1}) \chi(\gamma_{0,i_1} \circ (v_{i_1}, v_{i_2}) \circ \gamma_{0,i_2}^{-1}) \cdots \chi(\gamma_{0,i_{N-1}} \circ (v_{i_{N-1}}, v_{i_N}) \circ \gamma_{0,i_N}^{-1}) \chi(\gamma_{0,i_N} \circ (v_{i_N}, v_0) \circ \gamma_{00}^{-1}) \\ &= \chi(\{v_0\}) \chi((v_0, v_{i_1}) \circ (v_{i_1}, v_{i_2}) \circ \cdots \circ (v_{i_N}, v_0)) \chi(\{v_0\})^{-1} = \chi(\omega). \end{aligned}$$

This calculation is clearly independent of the choices of oriented paths $\{\tilde{\gamma}_{0i}\}$. Thus Hol maps $[\rho]$ exactly to χ , an element of $\text{Hom}(\pi_1(\Gamma, v_0), G)$; the independence of $\text{Hol}([\rho])$ as an element of $\text{Hom}(\pi_1(\Gamma), G)/G$ with respect to the choice of the base point v_0 follows from an essentially identical argument as given in the proof of Lemma 2.3. The last statement follows from Lemma 2.1. \square

Theorem 2.1 is closely related to classification theorems of flat connections and principal bundles with disconnected structure group (see, e.g. [Tau11, §13.6]) and [Ste51, §13.9]). The synchronizability of an edge potential ρ on connected graph Γ , which is equivalent to the triviality of \mathcal{B}_ρ (c.f. Proposition 1.2), can now be interpreted as the corresponding conjugacy class of $\text{Hom}(\pi_1(\Gamma), G)$. In fact, the conjugacy class corresponding to trivial bundles \mathcal{B}_ρ is also trivial and reflects the triviality of the holonomy of Γ . The proof of Corollary 2.2 further develops this observation.

Corollary 2.2. *For a connected graph Γ and topological group G , an edge potential $\rho \in C^1(\Gamma; G)$ is synchronizable if and only if $\text{Hol}_\rho(\Gamma)$ is trivial.*

Proof. Note that $\rho \in C^1(\Gamma; G)$ is synchronizable (see (1)) if and only if there exists $f \in C^0(\Gamma; G)$ such that

$$f_i^{-1} \rho_{ij} f_j = e \in G \quad \forall (i, j) \in E,$$

where e is the identity element of the structure group G . This is equivalent to saying that

$$[\rho] = [\epsilon] \in C^1(\Gamma; G)/C^0(\Gamma; G), \quad \text{where } \epsilon \in C^1(\Gamma; G) \text{ is defined as } \epsilon_{ij} = e \in G \text{ for all } (i, j) \in E, \quad (19)$$

which by Theorem 2.1 implies

$$\text{Hol}([\rho]) = \text{Hol}([\epsilon]) = \text{Id}_e \in \text{Hom}(\pi_1(\Gamma), G)/G,$$

where $\text{Id}_e : \pi_1(\Gamma) \rightarrow G$ is the constant map sending all oriented loops in Γ to the identity element $e \in G$. The conclusion follows immediately by noting that

$$\begin{aligned} \text{Hol}([\rho]) = \text{Id}_e &\Leftrightarrow \text{for any } v \in V, \text{ hol}_\rho(\omega) = e \text{ for all oriented loops } \omega \text{ based at } v \\ &\Leftrightarrow \text{Hol}_\rho(v) = \{e\} \text{ for all vertices } v \in V \\ &\Leftrightarrow \text{Hol}_\rho(\Gamma) \text{ is trivial} \end{aligned}$$

where for the last equivalence we invoked Lemma 2.4. \square

Corollary 2.2 on its own can be derived from an elementary argument. In fact, without descending hol_ρ from Ω_v to $\pi_1(\Gamma, v)$, we can still define $\text{Hol}_\rho(v)$ as the image $\text{hol}_\rho(\Omega_v)$, though hol_ρ is not injective as a group homomorphism from Ω_v to G . The triviality of $\text{Hol}_\rho(v)$ still implies the existence of a vertex potential $f \in C^0(\Gamma; G)$ compatible with $\rho \in C^1(\Gamma; G)$ (simply by setting $f_i = e$ on an arbitrarily chosen $v_i \in V$ and progressively propagating values of f to neighboring vertices), and *vice versa*. The exposition in this section, centered around Theorem 2.1, extends beyond this elementary argument and strives to unveil a complete geometric picture underlying the ‘‘correspondence between trivialities’’ discussed in Corollary 2.2. In future work we intend to pursue novel synchronization algorithms based on metric and symplectic structures on the moduli space of flat bundles (see, e.g. [AB83, Wei95, Hit90]).

2.2 A Twisted Hodge Theory for Synchronization Problems

In this section we relate synchronization to the first cohomology of a de Rham cochain complex on Γ with coefficients twisted by a representation space F of the structure group G . This can be interpreted as an instance of the standard de Rham cohomology of flat bundles (see e.g. [Xia12, GRR15, Sim94a, Sim94b]). The fibre bundles considered in this section are vector bundles (with fibre type F) associated with the principal bundle studied in Section 2.1. When the vector space F is equipped with a metric, the vector bundle inherits a compatible metric, with which a twisted Hodge Laplacian can be constructed; special cases of this twisted Laplacian in the lowest degree include the connection Laplacian [SW12, BSS13]. In this setting, synchronizability is realized as a condition on the dimension of the null space of the lowest degree twisted Hodge Laplacian, this is reminiscent of the classical Riemann-Hilbert correspondence between flat connections and locally constant sheaves. The spectral information of the twisted Hodge Laplacian serves as a quantification of the level of obstruction to synchronizability.

2.2.1 Flat Associated Bundles and Twisted Zero-Forms

Let \mathcal{B}_ρ be the synchronization principal bundle on Γ associated with $\rho \in C^1(\Gamma; G)$, as in Proposition 1.2, and F be a topological space on which G acts on the left as a topological transformation group. Denote the action of G on F as $\tau : G \rightarrow \text{Aut}(F)$. Consider the right action of G on $\mathcal{B}_\rho \times F$ as

$$(p, v) \mapsto (pg, \tau(g^{-1})v).$$

The orbit space of this action, conventionally denoted as $\mathcal{B}_\rho \times_G F$ or $\mathcal{B}_\rho[F]$, is referred to as the *F-bundle associated with principal bundle \mathcal{B}_ρ* , or *associated F-bundle* for short. We will denote the bundle projection as $\pi : \mathcal{B}_\rho[F] \rightarrow \Gamma$, and denote $\mathcal{B}_\rho[F]_x := \pi^{-1}(x)$ for the *fibre over $x \in \Gamma$* . Strictly speaking, the graph Γ should be distinguished from its underlying topological space, but we use the same notation Γ for both as long as the meaning is clear from the context.

The same open cover \mathcal{U} of Γ that trivializes \mathcal{B}_ρ also trivializes $\mathcal{B}_\rho[F]$. In fact, the bundle transition function of $\mathcal{B}_\rho[F]$ on any nonempty $U_i \cap U_j$ is the constant map $U_i \cap U_j \rightarrow \tau(\rho_{ij}) \in \text{Aut}(F)$, where $U_i \cap U_j \rightarrow \rho_{ij}$ is the constant bundle transition function from U_i to U_j for \mathcal{B}_ρ . Consequently, the associated bundle $\mathcal{B}_\rho[F]$ is also flat. Unless confusions arise, we shall refer to $\mathcal{B}_\rho[F]$ as the *flat associated F-bundle of \mathcal{B}_ρ* , and denote the local trivialization of the associated bundle over $U_i \in \mathcal{U}$ using the same notation $\phi_i : U_i \times F \rightarrow \mathcal{B}_\rho[F]$ as for the principal bundle \mathcal{B}_ρ .

In the context of synchronization problems, the most relevant associated bundles are those with fibre F being a vector space and structure group G being the general linear group $\text{GL}(F)$. These types of fibre bundles are commonly referred to as *vector bundles*. We will focus on flat associated vector bundles for the rest of the section, though the definition of *fibre projections* and *sections* extend literally to general fibre bundles. For simplicity of presentation, we will omit the notation τ and write the bundle transition functions again as ρ_{ij} (instead of $\tau(\rho_{ij})$), since its action on a vector space F is simply matrix-vector multiplication.

We now focus on *sections*, the analog of “functions” on smooth manifolds but with values in fibre bundles. For a general fibre bundle $\mathcal{E} \rightarrow \mathcal{B}$, a *local section* $s : U \rightarrow \mathcal{E}|_U$ of \mathcal{E} on an open set U of the base space \mathcal{B} is a continuous map from U to $\mathcal{E}|_U$ such that $\pi \circ s$ is identified on U . A *global section* of \mathcal{E} is a local section defined on the entire base space \mathcal{B} . We shall encode the data of synchronization problems into the language of sections of flat associated bundles. The discrete nature of the problem naturally motivates us to consider special classes of local and global sections that are “constant” within each open set in \mathcal{U} , in a sense to be made clear soon in local coordinates. The following notion of *fibre projection* is introduced to simplify notations involving local coordinates.

Definition 2.1. For any $i \in V$, define the *fibre projection* over $U_i \in \mathcal{U}$, denoted as

$$p_i : \mathcal{B}_\rho[F]|_{U_i} = \pi^{-1}(U_i) \longrightarrow F, \quad (20)$$

as the composition of $\phi_i^{-1} : \mathcal{B}_\rho[F]|_{U_i} \rightarrow U_i \times F$ with the canonical projection $U_i \times F \rightarrow F$. For any $x \in U_i$, the restriction of p_i to the fibre $\mathcal{B}_\rho[F]_x$, denoted as $p_{i,x} : \mathcal{B}_\rho[F]_x \rightarrow F$, is (by definition) simultaneously a homeomorphism between topological spaces and an isomorphism between vector spaces.

Definition 2.2 (Constant Local Sections). A *constant local section* $s : U_i \rightarrow \mathcal{B}_\rho[F]|_{U_i}$ of the bundle $\mathcal{B}_\rho[F]$ on open set $U_i \in \mathcal{U}$ is a local section of $\mathcal{B}_\rho[F]$ such that $p_{i,x}(s(x))$ is a constant element of F for all $x \in U_i$. We refer to the linear space of all constant local sections on $U_i \in \mathcal{U}$ as *constant twisted local 0-forms on U_i* , denoted as $\Omega_i^0(\Gamma; \mathcal{B}_\rho[F])$.

Clearly, a constant local section $s \in \Omega_i^0(\Gamma; \mathcal{B}_\rho[F])$ is unambiguously determined by evaluating s at vertex i , or equivalently by reading off the fibre projection image $s_i := p_i(s(i))$. We denote this characterization of s as

$$p_i(s(x)) \equiv s_i, \quad \forall x \in U_i, \quad s \in \Omega_i^0(\Gamma; \mathcal{B}_\rho[F]). \quad (21)$$

When we consider $x \in U_i \cap U_j$ where $U_i \cap U_j \neq \emptyset$ (i.e. when $(i, j) \in E$), it will be convenient to note that the fibre projection p_j evaluates $s(x)$ to $p_j(s(x)) = p_j \circ p_i^{-1}(p_i(s(x))) \equiv \rho_{ji}s_i$. This can be understood as a “change-of-coordinates” formula for constant local sections.

Let $C^0(\Gamma; F) := \{f : V \rightarrow F\}$ denote the linear space of F -valued 0-cochains on graph Γ . Every element f of $C^0(\Gamma; F)$ defines a collection of constant local sections $\{f^{(i)} : U_i \rightarrow \pi^{-1}(U_i) \mid U_i \in \mathcal{U}\}$, one for each $U_i \in \mathcal{U}$ with

$$f^{(i)}(x) := p_{i,x}^{-1}(f_i), \quad \forall x \in U_i. \quad (22)$$

We thus have the canonical identification

$$C^0(\Gamma; F) = \prod_{i \in V} \Omega_i^0(\Gamma; \mathcal{B}_\rho[F]). \quad (23)$$

Of course, the constant local sections $\{f^{(i)}\}$ specified by $f \in C^0(\Gamma; F)$ generally do not give rise to a global section of the bundle $\mathcal{B}_\rho[F]$, unless they “patch together” seamlessly on every nonempty intersection $U_i \cap U_j$, satisfying the condition

$$p_{i,x}^{-1}(f_i) = f^{(i)}(x) = f^{(j)}(x) = p_{j,x}^{-1}(f_j), \quad \forall x \in U_i \cap U_j \quad \Leftrightarrow \quad f_i = p_{i,x} \circ p_{j,x}^{-1}(f_j) = \rho_{ij}f_j, \quad \forall x \in U_i \cap U_j. \quad (24)$$

The right hand side of (24) is recognized as a solution to the synchronization problem with prescribed edge potential ρ . We have thus proved the following Lemma.

Lemma 2.5. *The constant local sections specified by $f \in C^0(\Gamma; F)$ define a global section on $\mathcal{B}_\rho[F]$ if and only if*

$$f_i = \rho_{ij}f_j, \quad \forall (i, j) \in E, \quad (25)$$

i.e., if and only if the vertex potential $f : V \rightarrow F$ is a solution to the F -synchronization problem over Γ with respect to the edge potential $\rho \in C^1(\Gamma; G)$.

When condition (25) is satisfied, the resulting global section constructed from constant local sections is special among all global sections of $\mathcal{B}_\rho[F]$ in that its restriction to each U_i has constant image under fibre projection over U_i . This type of global section will be of major interest in the remainder of this section.

Definition 2.3 (Locally Constant Global Section). A global section $s : \Gamma \rightarrow \mathcal{B}_\rho[F]$ is said to be *locally constant* if

$$p_i(s(x)) \equiv \text{const.} \quad \forall x \in U_i. \quad (26)$$

The linear space of all locally constant global sections on $\mathcal{B}_\rho[F]$ will be called *locally constant twisted global 0-forms* on Γ , denoted as $\Omega^0(\Gamma; \mathcal{B}_\rho[F])$.

Naturally, $\Omega^0(\Gamma; \mathcal{B}_\rho[F])$ embeds into $C^0(\Gamma; F)$ by

$$\begin{aligned} \Omega^0(\Gamma; \mathcal{B}_\rho[F]) &\hookrightarrow \prod_{i \in V} \Omega_i^0(\Gamma; \mathcal{B}_\rho[F]) = C^0(\Gamma; F) \\ s &\longmapsto (s|_{U_1}, \dots, s|_{U_n}) \end{aligned} \quad (27)$$

where $n = |V|$ stands for the total number of vertices in Γ . The objective of a F -synchronization problem over Γ with respect to $\rho \in C^1(\Gamma; G)$ can be interpreted in this geometric framework as searching for an element of $\Omega^0(\Gamma; \mathcal{B}_\rho[F])$ in the feasible domain $C^0(\Gamma; F)$.

The existence of global sections is crucial information for the structure of a fibre bundle. For principal bundles \mathcal{B}_ρ considered in Section 2.1, a single global section dictates the triviality of the bundle. Though the triviality of a principal bundle is equivalent to its associated vector bundle (see Proposition 2.1), $\mathcal{B}_\rho[F]$ is trivial if and only if it admits $d = \dim F$ global sections s^1, \dots, s^d that are *linearly independent* in the sense that s_x^1, \dots, s_x^d on each fibre

F_x are linearly independent as vectors in F (c.f. [MS74, Theorem 2.2]). A collection of linearly independent global sections are said to form a *global frame* (see e.g. [Lee03, Chapter 5]) for the vector bundle, since they define a basis (frame) for each fibre. As will be established in Proposition 2.1, the fact that the bundle $\mathcal{B}_\rho[F]$ is flat further reduces its triviality to finding d linearly independent locally constant global sections, for which linear independence only needs to be checked at the vertices of Γ . More precisely, adopting notation $s_i := p_i(s(i))$, $\forall i \in V$ for an arbitrary section s of $\mathcal{B}_\rho[F]$, we define the linear independence of locally constant global sections of $\mathcal{B}_\rho[F]$ as follows.

Definition 2.4. A collection of k ($1 \leq k \leq d = \dim F$) locally constant global sections $s^1, \dots, s^k \in \Omega^0(\Gamma; \mathcal{B}_\rho[F])$ are said to be *linearly independent* if s_i^1, \dots, s_i^k are linearly independent as vectors in F at every vertex $i \in V$.

By the embedding (27), any $s \in \Omega^0(\Gamma; \mathcal{B}_\rho[F])$ can be equivalently encoded into a vector of dimension nd , where $d = \dim F$ and $n = |V|$ stands for the number of vertices of Γ . In fact, just as for any vertex potential in $C^0(\Gamma; F)$, one simply needs to vertically stack the column vectors $\{s_i = p_i s(i) \mid i \in V\}$, the fibre projection images of s at each vertex. We shall write $[s]$ for such a vector of length nd that encodes $s \in \Omega^0(\Gamma; \mathcal{B}_\rho[F])$, and refer to the vector as the *representative vector* of the section. The linear independence of locally constant global sections is easily seen to be equivalent to the linear independence of the representative vectors of length nd , as the following Lemma clarifies.

Lemma 2.6. A collection of k ($1 \leq k \leq d = \dim F$) locally constant global sections s^1, \dots, s^k are linearly independent if and only if $[s^1], \dots, [s^k]$ are linear independent as vectors of length nd .

Proof. Since $\mathcal{B}_\rho[F]$ is a flat bundle and the graph Γ is assumed connected, the linear independence of locally constant global sections s^1, \dots, s^k is equivalent to the linear independence of vectors s_i^1, \dots, s_i^k at any vertex i : since $\rho_{ji} \in G$ are all invertible, vectors s_i^1, \dots, s_i^k are linearly independent if and only if $s_j^1 = \rho_{ji}s_i^1, \dots, s_j^k = \rho_{ji}s_i^k$ are linearly independent. For definiteness, let us fix $i = 1$. Write S for the nd -by- k matrix with $[s^j]$ as its j th column, S_1 for the $d \times k$ matrix with s_1^j as its j th column, $e = [1, \dots, 1]^\top$ for the column vector of length d with all entries equal to one, and P for the nd -by- nd block diagonal matrix with ρ_{j1} at its j th diagonal block (adopting the convention $\rho_{11} = I_{n \times n}$). The conclusion follows from the matrix identity

$$S = P(e \otimes S_1)$$

and

$$\text{rank}(S) = \text{rank}([1, \dots, 1]^\top) \cdot \text{rank}(S_1) = \text{rank}(S_1).$$

□

Remark 2.2. Note that the equivalence of two notions of linear independence only holds if we already know that s^1, \dots, s^k are global sections. For general $f^1, \dots, f^k \in C^0(\Gamma; F)$ that are linearly independent as nd -vectors, their corresponding representatives in $\prod_{i \in V} \Omega_i^0(\Gamma; \mathcal{B}_\rho[F])$ do not necessarily define global sections, nor are they in general linearly independent as constant local sections on each U_i . A simple example is to consider a graph Γ consisting of two vertices $V = \{v_1, v_2\}$ and only one edge connecting them, $F = \mathbb{R}^2$, and G is the trivial group consisting of only the 2×2 identity matrix: vectors $[f^1] = (1, 0, 1, 0)^\top$ and $[f^2] = (1, 0, 0, 1)^\top$ are linearly independent as vectors in \mathbb{R}^4 but do not define linearly independent constant local sections on U_1 .

With all essential concepts presented, we are ready to establish our main observation in this subsection.

Proposition 2.1. Let G be a topological transformation group acting on a (real or complex) d -dimensional vector space F on the left, $\Gamma = (V, E)$ be a connected undirected graph, and $\rho \in C^1(\Gamma; G)$ a G -valued edge potential. The following statements are equivalent:

- (i) \mathcal{B}_ρ is trivial;
- (ii) \mathcal{B}_ρ admits a global section;
- (iii) $\mathcal{B}_\rho[F]$ is trivial;
- (iv) $\mathcal{B}_\rho[F]$ admits $d = \dim F$ linearly independent locally constant global sections.

Proof. The equivalence of (i) and (iii) follows from [Ste51, Theorem 4.3]. The equivalence (i) \Leftrightarrow (ii) follows from standard differential geometry, see e.g. [Ste51, MS74]; similarly standard is the equivalence between (iii) and the existence of n linearly independent global sections on $\mathcal{B}_\rho[F]$. To show the equivalence (iv) \Leftrightarrow (iii), it suffices to prove that a trivial flat vector bundle $\mathcal{B}_\rho[F]$ admits d linearly independent global sections that are also locally constant. To see this, recall from Proposition 1.2 and Corollary 2.2 that

$$\begin{aligned}\mathcal{B}_\rho[F] \text{ is trivial} &\Leftrightarrow \mathcal{B}_\rho \text{ is trivial} \Leftrightarrow \text{Hol}([\rho]) \text{ is trivial} \\ &\Leftrightarrow \exists g \in \Omega_\rho^0(\Gamma; G) \text{ s.t. } g_i^{-1} \rho_{ij} g_j = e \ \forall (i, j) \in E \\ &\Leftrightarrow \exists g \in \Omega_\rho^0(\Gamma; G) \text{ s.t. } \rho_{ij} = g_i g_j^{-1} \ \forall (i, j) \in E.\end{aligned}$$

Let $\{e_1, \dots, e_d\}$ be a basis for F , and for each $k = 1, \dots, d$ define $s^k : \Gamma \rightarrow \mathcal{B}_\rho[F]$ as

$$s^k(x) = p_{i,x}^{-1}(g_i e_k) \quad \forall x \in U_i.$$

It is straightforward to verify by definition that s^k is a well-defined global section and is locally constant. That s^1, \dots, s^d are linearly independent as global sections follows from the fact that $p_{i,x} : F \rightarrow \mathcal{B}_\rho[F]_x$ are isomorphisms between vector spaces. \square

Remark 2.3. The global section in (ii) is also “locally constant” in a sense analogous to Definition 2.3 but for principal bundles; we do not introduce this definition here since global sections on principal bundles will not be pursued directly in this work.

Proposition 2.1 points out an alternative approach for determining the synchronizability of an edge potential $\rho \in C^1(\Gamma; G)$, at least when G is a matrix group $\text{GL}(F)$: it suffices to check the existence of $d = \dim F$ linearly independent locally constant global sections on the flat associated vector bundle $\mathcal{B}_\rho[F]$. Such existence can be stated as a cohomological obstruction. We will pursue such a formulation in the next section. In Section 2.2.3 we will utilize the inner product structure on F to reduce the structure group of a $\text{GL}(F)$ -bundle to $O(d)$ or $U(d)$, as commonly seen in synchronization problems [BKS16, BSS13, CKS15]. If the underlying fibre bundle is orientable, the same procedures further reduce the structure group to $\text{SO}(d)$ or $\text{SU}(d)$, corresponding to synchronization problems considered in [SZSH11, WS13].

2.2.2 Twisted One-Forms and De Rham Cohomology

In a smooth category, sections on a fibre bundle can be differentiated by a *covariant derivative*. The resulting object is a skew-symmetric “direction-dependent” section on the same bundle, or equivalently a section of a new fibre bundle which is the tensor product of the original fibre bundle with the bundle of 1-forms on the base manifold. We shall generalize this picture to the discrete/combinatorial setting for flat associated bundles $\mathcal{B}_\rho[F]$ that naturally arise in synchronization problems.

Recall from discrete Hodge theory [CdGDS13, Lim15, DHLM05, JLYY11, JG13] that discrete 0-forms and 1-forms on a graph Γ are defined as

$$\Omega^0(\Gamma) := \{f : V \rightarrow \mathbb{K}\}, \quad \Omega^1(\Gamma) := \{\omega : E \rightarrow \mathbb{K} \mid \omega_{ij} = -\omega_{ji} \ \forall (i, j) \in E\},$$

where $\mathbb{K} = \mathbb{C}$ or \mathbb{R} , which we will assume to be the number field for the vector space F . Let us define a local version of $\Omega^1(\Gamma)$ by

$$\Omega_i^1(\Gamma) := \{\omega : N_i \rightarrow \mathbb{K} \mid \omega_{jk} = -\omega_{kj} \ \forall (j, k) \in N_i\}, \quad \text{where } N_i := \{(j, k) \in E \mid j = i \text{ or } k = i\}. \quad (28)$$

In other words, elements of $\Omega_i^1(\Gamma)$ are restrictions of elements of $\Omega^1(\Gamma)$ to U_i . By a partition of unity argument, it is straightforward to identify $\Omega^1(\Gamma)$ with

$$\left\{ \left(\omega^{(1)}, \dots, \omega^{(n)} \right) \in \prod_{i \in V} \Omega_i^1(\Gamma) \mid \omega_{ij}^{(i)} = \omega_{ij}^{(j)} (= -\omega_{ji}^{(j)} = -\omega_{ji}^{(i)}) \right\}. \quad (29)$$

Definition 2.5 (Constant Local 1-forms). A *constant twisted local 1-form* on open set $U_i \in \mathcal{U}$ is a local section of $\Omega_i^0(\Gamma; \mathcal{B}_\rho[F]) \otimes \Omega_i^1(\Gamma)$. Equivalently, a constant twisted local 1-form on U_i is a map $\omega : U_i \times N_i \rightarrow \mathcal{B}_\rho[F]$ such that:

(i) For any $(j, k) \in N_i$, $\omega_{jk} : U_i \rightarrow \mathcal{B}_\rho[F]$ is a constant local section on U_i , i.e.

$$p_{i,x}(\omega_{ij}(x)) \equiv \text{const.};$$

(ii) For any $x \in U_i$ and any $U_j \in \mathfrak{U}$ such that $U_i \cap U_j \neq \emptyset$, $\omega_{ij}(x) = -\omega_{ji}(x)$.

We denote the linear space of all constant twisted local 1-forms on U_i as $\Omega_i^1(\Gamma; \mathcal{B}_\rho[F])$.

A similar notion of globally defined twisted 1-forms will also be of interest. In the discrete setting of synchronization problems, it suffices to consider twisted global 1-forms that are locally constant under fibre projections.

Definition 2.6 (Locally Constant Global 1-forms). A *locally constant twisted global 1-form* is a section of the tensor product bundle $\Omega^0(\Gamma; \mathcal{B}_\rho[F]) \otimes \Omega^1(\Gamma)$. In other words, a locally constant twisted global 1-form is a map

$$\omega : \{(i, (j, k)) \mid i \in V, (j, k) \in N_i\} \rightarrow \mathcal{B}_\rho[F]$$

such that:

(i) For any $U_i \in \mathfrak{U}$, $\omega|_{U_i}$ is a constant twisted local 1-form on U_i ;

(ii) For any $x \in U_i \cap U_j \neq \emptyset$, $p_{i,x}(\omega|_{U_i}(x)) = \rho_{ij} p_{j,x}(\omega|_{U_j}(x))$.

We denote the linear space of all locally constant twisted global 1-forms on Γ as $\Omega^1(\Gamma; \mathcal{B}_\rho[F])$.

The definition of $\Omega^1(\Gamma; \mathcal{B}_\rho[F])$ already characterized the condition under which a given collection of constant twisted local 1-forms $\{\omega^{(i)} \in \Omega_i^1(\Gamma; \mathcal{B}_\rho[F])\}$, one for each $U_i \in \mathfrak{U}$, can be patched to form a locally constant twisted global 1-form. Similar to (24), it suffices to check the compatibility under “change of coordinates”.

Lemma 2.7. A collection of constant twisted local 1-forms $\{\omega^{(i)} \in \Omega_i^1(\Gamma; \mathcal{B}_\rho[F])\}$ defines a locally constant twisted global 1-form if and only if

$$p_{i,x}(\omega_{ij}^{(i)}(x)) = \rho_{ij} p_{j,x}(\omega_{ij}^{(j)}(x)), \quad \forall (i, j) \in E. \quad (30)$$

Since both sides of equality (30) are constants, we shall simplify (30) as

$$p_i(\omega_{ij}^{(i)}) = \rho_{ij} p_j(\omega_{ij}^{(j)}), \quad \forall (i, j) \in E. \quad (31)$$

A significant difference between $\Omega^1(\Gamma; \mathcal{B}_\rho[F])$ and $\Omega^0(\Gamma; \mathcal{B}_\rho[F])$ is that a locally constant twisted global 1-form does not naturally arise from an F -valued 1-cochain in $C^1(\Gamma; F) := \{\omega : E \rightarrow F \mid \omega_{ij} = -\omega_{ji}\}$, and these cochains play an essential role in the discrete Hodge theory. For instance, if we set $\omega_{ij}^{(i)}(x) := p_{i,x}^{-1}(\omega_{ij}) =: -\omega_{ji}^{(i)}(x)$ for all $x \in U_i$, then $\{\omega^{(i)} \mid i \in V\}$ gives rise to a twisted global 1-form if and only if

$$\omega_{ij} = p_{i,x}(\omega_{ij}^{(i)}(x)) = \rho_{ij} p_{j,x}(\omega_{ij}^{(j)}(x)) = -\rho_{ij} p_{j,x}(\omega_{ji}^{(j)}(x)) = -\rho_{ij} \omega_{ji}, \quad \forall (i, j) \in E,$$

a condition that is generally not satisfied unless $\rho_{ij} \equiv e \in G$ for all $(i, j) \in E$. This observation indicates that $C^1(\Gamma; F)$ is not a geometric object naturally associated with the structure of the vector bundle $\mathcal{B}_\rho[F]$, but rather a special case of $\Omega^1(\Gamma; \mathcal{B}_\rho[F])$ when the vector bundle $\mathcal{B}_\rho[F]$ is trivial. In this case $\rho_{ij} = g_i g_j^{-1}$, $\forall (i, j) \in E$ for a G -valued vertex potential $g : V \rightarrow G$ and the “gauge-transformed” constant twisted local 1-forms $\{p_{i,x}^{-1}(g_i \omega_{ij}) \mid i \in V, (i, j) \in E\}$ satisfy the compatibility condition (30). Exploring the action of the gauge group on twisted forms will be considered in the future.

With an appropriate notion of twisted 1-forms, we are ready to define the twisted differential operator on twisted 0-forms. This operation is a discrete analog of the covariant derivatives for smooth fibre bundles, and in the meanwhile, a fibre bundle analog of the discrete exterior derivative in discrete Hodge theory.

Definition 2.7 (Twisted Differential on Twisted 0-Cochains). For $\rho \in C^1(\Gamma; G)$ and $U_i \in \mathfrak{U}$, the ρ -twisted differential is a linear operator taking any $f \in C^0(\Gamma; F) = \prod_{i \in V} \Omega_i^0(\Gamma; \mathcal{B}_\rho[F])$ to a collection of n constant twisted local 1-forms, one for each $U_i \in \mathfrak{U}$

$$d_\rho : \prod_{i \in V} \Omega_i^0(\Gamma; \mathcal{B}_\rho[F]) \longrightarrow \prod_{i \in V} \Omega_i^1(\Gamma; \mathcal{B}_\rho[F])$$

$$f \longmapsto \left((d_\rho f)^{(1)}, \dots, (d_\rho f)^{(n)} \right)$$

where each $(d_\rho f)^{(i)} \in \Omega_i^1(\Gamma; \mathcal{B}_\rho[F])$ is defined as

$$(d_\rho f)_{ij}^{(i)}(x) := p_{i,x}^{-1}(f_i - \rho_{ij}f_j) =: -(d_\rho f)_{ji}^{(i)}(x), \quad \forall U_i \in \mathfrak{U}, \forall x \in U_i \cap U_j \neq \emptyset, f \in \Omega_\rho^0(\Gamma; F). \quad (32)$$

Though d_ρ is defined as a linear operator mapping into a collection of constant twisted local 1-forms, a somewhat surprising fact is that these constant twisted local 1-forms do patch together to form an element of $\Omega^1(\Gamma; \mathcal{B}_\rho[F])$.

Proposition 2.2. *The twisted differential d_ρ maps $C^0(\Gamma; F)$ into $\Omega^1(\Gamma; \mathcal{B}_\rho[F])$.*

Proof. It suffices to check (31) for the collection of constant twisted local 1-forms $\left\{ (d_\rho f)^{(i)} \mid i = 1, \dots, n \right\}$. In fact,

$$p_i \left((d_\rho f)_{ij}^{(i)} \right) = f_i - \rho_{ij}f_j = -\rho_{ij}(f_j - \rho_{ji}f_i) = -\rho_{ij}p_j \left((d_\rho f)_{ji}^{(j)} \right) = \rho_{ij}p_j \left((d_\rho f)_{ij}^{(j)} \right) \quad \forall (i, j) \in E.$$

□

Since the graph Γ (viewed as a simplicial complex) does not contain any 2-simplices, d_ρ is the only differential needed for specifying the ρ -twisted chain complex

$$0 \longrightarrow C^0(\Gamma; F) \xrightarrow{d_\rho} \Omega^1(\Gamma; \mathcal{B}_\rho[F]) \longrightarrow 0. \quad (33)$$

The only non-trivial cohomology group in this de Rham-type chain complex is at the 0-th order

$$H_\rho^0(\Gamma; \mathcal{B}_\rho[F]) := \ker d_\rho.$$

Proposition 2.3. $\ker d_\rho = \Omega^0(\Gamma; \mathcal{B}_\rho[F])$.

Proof. Note in the definition (32) that

$$f \in \ker d_\rho \quad \Leftrightarrow \quad f_i = \rho_{ij}f_j, \quad \forall (i, j) \in E.$$

The conclusion then follows from Lemma 2.5. □

By Proposition 2.1, detecting the synchronizability of $\rho \in C^1(\Gamma; G)$ now reduces to checking if $\dim \ker d_\rho = \dim F$ holds. Furthermore, in scenarios where this dimension equality does not hold, $\dim \ker d_\rho$ still provides a quantitative measure for the extent to which synchronizability fails. In this sense, the cohomology group $H_\rho^0(\Gamma; F)$ serves as the opposite of a “topological obstruction” to the synchronizability of $\rho \in C^1(\Gamma; G)$.

2.2.3 Twisted Hodge Theory and Synchronizability

In the remainder of this section, we will focus on flat associated bundles $\mathcal{B}_\rho[F]$ with the vector space F equipped with an inner product $\langle \cdot, \cdot \rangle_F : F \times F \rightarrow \mathbb{K}$, where $\mathbb{K} = \mathbb{C}$ or \mathbb{R} depending on the vector space F . This inner product on F will be further assumed with G -invariance, in the sense that

$$\langle gx, gy \rangle_F = \langle x, y \rangle_F \quad \forall x, y \in F, g \in G.$$

In the terminology of representation theory, we assume that the representation of G on F is *unitary* (c.f. [BtD03, §II.1]). This inner product introduces other related concepts into the geometric framework:

- F is equipped with a G -invariant norm $\|x\|_F = \langle x, x \rangle_F$ for all $x \in F$, which further induces an operator norm on G via duality

$$\|g\| := \sup_{\substack{m \in F \\ \|m\| \neq 0}} \frac{\|gm\|}{\|m\|}, \quad \forall g \in G.$$

For simplicity, we will use the same notation for the norm on F and the dual norm on G .

- For any $g \in G$, its *formal adjoint* with respect to the inner product $\langle \cdot, \cdot \rangle_F$, denoted as g^* , is defined as

$$\langle gx, y \rangle = \langle x, g^*y \rangle \quad \forall x, y \in F.$$

Note that $\|g^*\| = \|g\|$ for any $g \in G$.

- The twisted 0-cochains $C^0(\Gamma; F)$ and locally constant twisted global 1-forms $\Omega^1(\Gamma; \mathcal{B}_\rho[F])$ are equipped with inner products and norms induced from the G -invariant inner product $\langle \cdot, \cdot \rangle_F$, as follows:

$$\langle f, g \rangle := \sum_{i \in V} d_i \langle f_i, g_i \rangle_F, \quad \forall f, g \in C^0(\Gamma; F) \quad (34)$$

$$\langle \omega, \eta \rangle := \frac{1}{2} \sum_{i \in V} \sum_{j: (i,j) \in E} w_{ij} \langle p_i(\omega_{ij}^{(i)}), p_i(\eta_{ij}^{(i)}) \rangle_F, \quad \forall \omega, \eta \in \Omega^1(\Gamma; \mathcal{B}_\rho[F]) \quad (35)$$

$$\|f\| := \langle f, f \rangle^{\frac{1}{2}}, \quad \|\omega\| := \langle \omega, \omega \rangle^{\frac{1}{2}}, \quad \forall f \in C^0(\Gamma; F), \omega \in \Omega^1(\Gamma; \mathcal{B}_\rho[F]) \quad (36)$$

where w_{ij} is the weight on edge $(i, j) \in E$ and $d_i = \sum_{j: (i,j) \in E} w_{ij}$ is the weighted degree of vertex v_i . Note that by the G -invariance of $\langle \cdot, \cdot \rangle_F$ the sum in (35) can be equivalently written as (see Appendix A for a quick calculation)

$$\langle \omega, \eta \rangle = \sum_{(i,j) \in E} w_{ij} \langle p_i(\omega_{ij}^{(i)}), p_i(\eta_{ij}^{(i)}) \rangle_F. \quad (37)$$

Through local trivializations, an inner product on the vector space F also induces a *bundle metric* on $\mathcal{B}_\rho[F]$, i.e., a section of the second symmetric power of the dual bundle of $\mathcal{B}_\rho[F]$ which restricts to each fibre as a symmetric positive definite quadratic form. As is well known (see e.g. [Tau11, Chapter 7]), a bundle metric can be used to reduce the structure group of a vector bundle from $GL(F)$ to $O(d)$ or $U(d)$, where $d = \dim(F)$. It suffices to consider global sections of $\mathcal{B}_\rho[F]$ for $\rho \in C^1(\Gamma; O(d))$ or $\rho \in C^1(\Gamma; U(d))$ for many synchronization problems of practical interest [BSS13, BKS16, CKS15], instead of requiring $\rho \in C^1(\Gamma; GL(d; \mathbb{R}))$ or $\rho \in C^1(\Gamma; GL(d; \mathbb{C}))$. Other important types of synchronization problem involving $SO(d)$ and $SU(d)$ can be treated in this geometric framework as determining global sections of *orientable* vector bundles (see e.g. [Tau11, Chapter 7] or [BT82, Proposition 6.4]). Also note that $\rho_{ij}^{-1} = \rho_{ij}^*$ for edge potentials in all these special matrix groups. Since the bundle reduction allows us to focus only on synchronization problems with orthogonal or unitary matrices, without loss of generality, we will always assume the edge potentials satisfy $\rho_{ji} = \rho_{ij}^{-1} = \rho_{ij}^*$ for all $(i, j) \in E$. The same is assumed in [BSS13, BKS16].

The inner product structures on $C^0(\Gamma; F)$ and $\Omega^1(\Gamma; \mathcal{B}_\rho[F])$ enable us to define the ρ -twisted codifferential $\delta_\rho : \Omega^1(\Gamma; \mathcal{B}_\rho[F]) \rightarrow C^0(\Gamma; F)$, the formal adjoint operator of the twisted differential $d_\rho : C^0(\Gamma; F) \rightarrow \Omega^1(\Gamma; \mathcal{B}_\rho[F])$ in the chain complex (33), eventually leading to a twisted Hodge theory for synchronization problems. The definition of δ_ρ is consistent with the discrete divergence operator in discrete Hodge theory [JLYY11, JG13]:

$$\begin{aligned} \delta_\rho : \Omega^1(\Gamma; \mathcal{B}_\rho[F]) &\longrightarrow C^0(\Gamma; F) \\ \theta &\longmapsto ((\delta_\rho \theta)|_{U_1}, \dots, (\delta_\rho \theta)|_{U_n}) \end{aligned} \quad (38)$$

where each $(\delta_\rho \theta)|_{U_i} \in F$ is defined by

$$(\delta_\rho \theta)|_{U_i}(x) = p_{i,x}^{-1} \left(\frac{1}{d_i} \sum_{j: (i,j) \in E} w_{ij} p_i(\theta_{ij}^{(i)}) \right) \quad \forall x \in U_i, \theta \in \Omega^1(\Gamma; \mathcal{B}_\rho[F]) \quad (39)$$

or equivalently

$$(\delta_\rho \theta)_i = p_i((\delta_\rho \theta)|_{U_i}(i)) = \frac{1}{d_i} \sum_{j: (i,j) \in E} w_{ij} p_i(\theta_{ij}^{(i)}) \quad \forall i \in V, \theta \in \Omega^1(\Gamma; \mathcal{B}_\rho[F]). \quad (40)$$

Proposition 2.4. *With respect to the inner products (34) and (35), the twisted codifferential $\delta_\rho : \Omega^1(\Gamma; \mathcal{B}_\rho[F]) \rightarrow C^0(\Gamma; F)$ defined by (39) is the formal adjoint of the twisted differential $d_\rho : C^0(\Gamma; F) \rightarrow \Omega^1(\Gamma; \mathcal{B}_\rho[F])$ defined by (32).*

Proof. Note that for any $f \in C^0(\Gamma; F)$, $\theta \in \Omega^1(\Gamma; \mathcal{B}_\rho[F])$,

$$\begin{aligned} \langle f, \delta_\rho \theta \rangle &= \sum_{i \in V} d_i \left\langle f_i, \frac{1}{d_i} \sum_{j: (i,j) \in E} w_{ij} p_i(\theta_{ij}^{(i)}) \right\rangle_F = \sum_{i \in V} \sum_{j: (i,j) \in E} \langle f_i, w_{ij} p_i(\theta_{ij}^{(i)}) \rangle_F \\ &= \sum_{(i,j) \in E} [\langle f_i, w_{ij} p_i(\theta_{ij}^{(i)}) \rangle_F + \langle f_j, w_{ji} p_j(\theta_{ji}^{(j)}) \rangle_F] \\ &= \sum_{(i,j) \in E} \langle f_i, w_{ij} p_i(\theta_{ij}^{(i)}) \rangle_F + \sum_{(i,j) \in E} \langle f_j, w_{ji} p_j(\theta_{ji}^{(j)}) \rangle_F =: (I) + (II). \end{aligned}$$

We keep the term (I) intact and manipulate term (II) using $w_{ij} = w_{ji}$ and the G -invariance of $\langle \cdot, \cdot \rangle_F$:

$$(II) = \sum_{(i,j) \in E} \langle \rho_{ij} f_j, w_{ji} \rho_{ij} p_j(\theta_{ji}^{(j)}) \rangle_F \stackrel{(*)}{=} \sum_{(i,j) \in E} \langle \rho_{ij} f_j, w_{ij} p_i(\theta_{ji}^{(i)}) \rangle_F \stackrel{(**)}{=} - \sum_{(i,j) \in E} \langle \rho_{ij} f_j, w_{ij} p_i(\theta_{ij}^{(i)}) \rangle_F,$$

where we used $\rho_{ij} p_j(\theta_{ji}^{(j)}) = p_i(\theta_{ji}^{(i)})$ (the compatibility condition (31)) at (*), and the skew-symmetry $\theta_{ji}^{(i)} = -\theta_{ij}^{(i)}$ at (**). Re-combining (I) and (II), we conclude that

$$\begin{aligned} \langle f, \delta_\rho \theta \rangle &= (I) + (II) = \sum_{(i,j) \in E} \langle f_i, w_{ij} p_i(\theta_{ij}^{(i)}) \rangle_F - \sum_{(i,j) \in E} \langle \rho_{ij} f_j, w_{ij} p_i(\theta_{ij}^{(i)}) \rangle_F \\ &= \sum_{(i,j) \in E} \langle f_i - \rho_{ij} f_j, w_{ij} p_i(\theta_{ij}^{(i)}) \rangle_F = \sum_{(i,j) \in E} w_{ij} \left\langle p_i((d_\rho f)_{ij}^{(i)}), p_i(\theta_{ij}^{(i)}) \right\rangle_F = \langle d_\rho f, \theta \rangle. \end{aligned}$$

□

The chain complex (33) is now also equipped with formal adjoints:

$$0 \rightleftharpoons C^0(\Gamma; F) \xrightleftharpoons[\delta_\rho]{d_\rho} \Omega^1(\Gamma; \mathcal{B}_\rho[F]) \rightleftharpoons 0. \quad (41)$$

Two *twisted Hodge Laplacians* can be constructed from this chain complex:

$$\Delta_\rho^{(0)} := \delta_\rho d_\rho : C^0(\Gamma; F) \longrightarrow C^0(\Gamma; F), \quad (42)$$

$$\Delta_\rho^{(1)} := d_\rho \delta_\rho : \Omega^1(\Gamma; \mathcal{B}_\rho[F]) \longrightarrow \Omega^1(\Gamma; \mathcal{B}_\rho[F]). \quad (43)$$

It is straightforward to see from these definitions that both twisted Laplacians are positive definite. In view of Hodge theory, it would be of interest to investigate the *harmonic forms* in the complex (41), the kernels of $\Delta_\rho^{(0)}$ and $\Delta_\rho^{(1)}$.

Lemma 2.8. $\ker d_\rho = \ker \Delta_\rho^{(0)}$ and $\ker \delta_\rho = \ker \Delta_\rho^{(1)}$.

Proof. Clearly $\ker d_\rho \subset \ker \Delta_\rho^{(0)}$. For the reverse inclusion, note that by adjointness

$$0 = \langle f, \Delta_\rho^{(0)} f \rangle = \|d_\rho f\|^2 \quad \forall f \in \ker \Delta_\rho^{(0)},$$

which implies $d_\rho f = 0$. The equality involving $\ker \Delta_\rho^{(1)}$ follows from a similar argument. □

The following decomposition results follow from standard Hodge-theoretic arguments.

Theorem 2.3. $C^0(\Gamma; F) = \ker \Delta_\rho^{(0)} \oplus \operatorname{im} \delta_\rho = \ker d_\rho \oplus \operatorname{im} \delta_\rho$, $\Omega^1(\Gamma; \mathcal{B}_\rho[F]) = \operatorname{im} d_\rho \oplus \ker \Delta_\rho^{(1)} = \operatorname{im} d_\rho \oplus \ker \delta_\rho$.

Proof. We only present the proof for the decomposition of $C^0(\Gamma; F)$; the decomposition for $\Omega^1(\Gamma; F)$ is similar. First note that both $C^0(\Gamma; F)$ and $\Omega^1(\Gamma; \mathcal{B}_\rho[F])$ are finite dimensional. The subspace $\ker d_\rho$ and $\text{im } \delta_\rho$ are orthogonal with respect to the inner product (34), since if $f \in \ker d_\rho$ and $\delta_\rho \theta \in \text{im } \delta_\rho$,

$$\langle f, \delta_\rho \theta \rangle = \langle d_\rho f, \theta \rangle = 0.$$

It remains to prove that each $f \in C^0(\Gamma; F)$ can be decomposed into a linear combination of elements in $\ker \Delta_\rho^{(0)}$ and $\text{im } \delta_\rho$. If $d_\rho f = 0$, the decomposition is trivial. Otherwise, consider the following Poisson equation:

$$\Delta_\rho^{(1)} \theta = d_\rho f. \quad (44)$$

We claim that equation (44) has a solution $\theta \in \Omega^1(\Gamma; \mathcal{B}_\rho[F])$ as long as $d_\rho f \neq 0$. In fact, by Fredholm alternative (see e.g. an exposition for the finite dimensional case in [Lim15] which suffices for our purpose), if $d_\rho f \notin \text{im } \Delta_\rho^{(1)}$, then $d_\rho f \in \ker \Delta_\rho^{(1)} = \ker \delta_\rho$; however, $d_\rho f \perp \ker \delta_\rho$ since

$$\langle d_\rho f, \omega \rangle = \langle f, \delta_\rho \omega \rangle = 0 \quad \forall \omega \in \ker \delta_\rho.$$

This proves that (44) has a solution $\theta \in \Omega^1(\Gamma; \mathcal{B}_\rho[F])$ for $d_\rho f \neq 0$. We can thus split $f \in C^0(\Gamma; F)$ into

$$f = (f - \delta_\rho \theta) + \delta_\rho \theta,$$

in which $\delta_\rho \theta \in \text{im } \delta_\rho$, and $f - \delta_\rho \theta \in \ker d_\rho$ since

$$d_\rho (f - \delta_\rho \theta) = d_\rho f - d_\rho \delta_\rho \theta = d_\rho f - \Delta_\rho^{(1)} \theta = 0.$$

□

Remark 2.4. Proposition 2.3 and Theorem 2.3 completely characterized the embedding (27): the orthogonality complement (with respect to the inner product (34)) of the linear space of solutions to the F -synchronization problem on Γ with respect to $\rho \in \Omega^1(\Gamma; G)$ is exactly the image of the twisted codifferential (39). In fact, one recognizes from $C^0(\Gamma; F) = \ker \Delta_\rho^{(0)} \oplus \text{im } \delta_\rho$, $\ker \Delta_\rho^{(0)} = \ker d_\rho$, and $H_\rho^0(\Gamma; \mathcal{B}_\rho[F]) = \ker d_\rho$ the well-known Hodge theorem $H_\rho^0(\Gamma; \mathcal{B}_\rho[F]) = \ker \Delta_\rho^{(0)}$.

2.2.4 Graph Connection Laplacian and Cheeger-Type Inequalities for Graph Frustration

In this section, we connect our geometric framework to the computational aspects of synchronization algorithms. As pointed out in Section 2.2.2, for cases where $G = O(d)$ or $G = U(d)$, the synchronizability of an edge potential $\rho \in C^1(\Gamma; G)$ is equivalent to whether or not the equality $\dim \ker d_\rho = d$ holds. Lemma 2.8 reduces the synchronizability further to the dimension of $\ker \Delta_\rho^{(0)}$. With identification (23), it can be noticed that $\Delta_\rho^{(0)}$ is exactly the *graph connection Laplacian* (GCL) in the literature of synchronization problems, random matrix theory, and manifold learning (see e.g. [SW12, BSS13, KLP⁺16, EW16]). Recall from [BSS13] that the graph connection Laplacian for graph Γ and edge potential $\rho \in C^1(\Gamma; G)$ is defined as

$$L_1 = D_1 - W_1 \quad (45)$$

where $W_1 \in \mathbb{K}^{nd \times nd}$ is a $n \times n$ block matrix with $w_{ij} \rho_{ij} \in \mathbb{K}^{d \times d}$ at its (i, j) th block, and $D_1 \in \mathbb{K}^{nd \times nd}$ is block diagonal with $d_i I_{d \times d} \in \mathbb{K}^{d \times d}$ at its (i, i) th block. To see that $\Delta_\rho^{(0)}$ coincides with L_1 , notice that

$$\langle f, \Delta_\rho^{(0)} f \rangle = \|d_\rho f\|^2 = \sum_{(i,j) \in E} w_{ij} \|f_i - \rho_{ij} f_j\|^2 = \frac{1}{2} \sum_{i,j \in V} w_{ij} \|f_i - \rho_{ij} f_j\|^2 = \frac{1}{2} [f]^\top L_1 [f], \quad \forall f \in C^0(\Gamma; F).$$

Theorem 2.3 translates into this combinatorial setting as a decomposition result for the matrix L_1 , as presented below in Proposition 2.5. We denote $n = |V|$ and $m = |E|$ for the graph $\Gamma = (V, E)$.

Proposition 2.5. *The graph connection Laplacian $L_1 \in \mathbb{K}^{nd \times nd}$ admits a decomposition*

$$L_1 = [\delta_\rho][d_\rho], \quad [\delta_\rho] \in \mathbb{K}^{nd \times md}, [d_\rho] \in \mathbb{K}^{md \times nd}, \quad (46)$$

where $[d_\rho]$ is an m -by- n block matrix in which the (i, j) th block is given by

$$[d_\rho]_{ij} = \begin{cases} I_{d \times d} & \text{if edge } i \text{ starts at vertex } j, \\ -w_{kj}\rho_{kj} & \text{if edge } i \text{ starts at vertex } k \text{ and ends at vertex } j, \\ 0 & \text{otherwise,} \end{cases} \quad (47)$$

and $[\delta_\rho]$ is an n -by- m block matrix in which the (i, j) th block is given by

$$[\delta_\rho]_{ij} = \begin{cases} \frac{w_{ik}}{d_i} I_{d \times d} & \text{if edge } j \text{ starts at vertex } i \text{ and ends at vertex } k, \\ 0 & \text{otherwise.} \end{cases} \quad (48)$$

Note that here each edge (i, j) appears twice in E with opposite orientations.

The Hodge decomposition (46) immediately leads to the following observation, which reflects the geometric fact that there do not exist more than $n = \dim F$ linearly independent global sections on the vector bundle $\mathcal{B}_\rho[F]$.

Proposition 2.6. *The dimension of the null eigenspace of L_1 can not exceed n , the dimension of both the column space of $[d_\rho]$ and the row space of $[\delta_\rho]$.*

By Lemma 2.6, if there are d linearly independent vectors in the kernel space of L_1 , then they give rise to d locally constant global sections on $\mathcal{B}_\rho[F]$ that are also linearly independent as global sections, which indicates the triviality of the vector bundle $\mathcal{B}_\rho[F]$ and the synchronizability of $\rho \in C^1(\Gamma; G)$. Note that an analogy of this result for graphs with multiple connected components also holds, though we assumed Γ is connected throughout this paper: a graph with $k \geq 1$ connected components and a prescribed $O(d)$ -valued edge potential is synchronizable if and only if the dimension of the null eigenspace of L_1 is kd . This geometric picture is consistent with the main spectral relaxation algorithm [BSS13, Algorithm 2.5] when the edge potential is synchronizable. Basically, the spectral relaxation procedure works as follows: first, extract d eigenvectors x^1, \dots, x^d corresponding to the smallest d eigenvalues of L_1 ; second, form the $nd \times d$ matrix $X = [x^1, \dots, x^d]$ and split it vertically into n blocks X_1, \dots, X_n of equal size $d \times d$; finally, find the closest orthogonal matrix O_i to each X_i by polar decomposition, and construct the desired synchronizing vertex potential $f \in C^0(\Gamma; O(d))$ by setting $f_i = O_i$. Since $\ker \Delta_\rho^{(0)} = d$ for any synchronizable edge potential ρ , the d eigenvectors x^1, \dots, x^d of L_1 all lie in the null eigenspace of L_1 , which provide exactly the d linearly independent global sections needed to trivialize the vector bundle; all that remains for obtaining a desired synchronizing vertex potential is to rescale the columns of each block $X_i \in \mathbb{K}^{d \times d}$ to achieve orthonormality, which is exactly what is done in the polar decomposition step when ρ is synchronizable. The twisted cohomology framework developed in this section suggests the following improvements when applying the spectral relaxation algorithm for determining synchronizability of a given edge potential:

- (1) Instead of checking $\dim \ker \Delta_\rho^{(0)}$, one can simply check $\dim \ker d_\rho$ or $\dim \ker \delta_\rho$ (which gives $\dim \operatorname{im} \delta_\rho$). Both $[d_\rho]$ and $[\delta_\rho]$ matrices are much smaller in size compared with L_1 , and the dimension can be determined by QR decomposition rather than the more expensive eigen-decomposition;
- (2) Instead of performing polar decomposition for each $d \times d$ block X_i , which involves the relatively more expensive SVD, it suffices to invoke a Gram-Schmidt orthonormalization. If synchronizability of ρ is confirmed by the dimension test in a previous step, the Gram-Schmidt procedure can be performed for the entire matrix $X \in \mathbb{K}^{nd \times d}$ in one pass (with a minor modification of keeping the columns to have norm n instead of 1), as opposed to being carried out for each individual block X_i .

Remark 2.5. *The Hodge decomposition (46) also suggests an alternative approach to obtaining n linearly independent locally constant global sections on $\mathcal{B}_\rho[F]$: instead of directly solving for the null eigenspace of L_1 , we can look for the orthogonal complement of $\operatorname{im} [d_\rho]$. Note, however, that the domain of $[d_\rho]$ should not be taken as the entire \mathbb{K}^{md} , since δ_ρ is defined on $\Omega^1(\Gamma; \mathcal{B}_\rho[F])$, in which elements satisfy the compatibility condition (31). Constructing such a basis matrix $B \in \mathbb{K}^{md \times md}$ and computing the orthogonal complement of the column space of $[d_\rho]B$ turns out not to be much simpler than finding the orthogonal complement of L_1 (i.e. finding the null eigenspace of L_1 directly).*

In the more general setting where the edge potential ρ is not assumed synchronizable, the geometric picture becomes much more involved. Of central importance to the relaxation algorithms and Cheeger inequalities in [BSS13] is to minimize the *frustration* of a graph Γ with respect to a prescribed group potential:

$$\begin{aligned} \nu(\Gamma) &= \inf_{g \in C^0(\Gamma; O(d))} \nu(g) \\ &= \inf_{g \in C^0(\Gamma; O(d))} \frac{1}{2d} \frac{1}{\text{vol}(\Gamma)} \sum_{i,j \in V} w_{ij} \|g_i - \rho_{ij} g_j\|_F^2, \quad \text{where } \text{vol}(\Gamma) = \sum_{i \in V} d_i. \end{aligned} \quad (49)$$

As shown in (the proof of) Proposition 2.1, an $O(d)$ -valued edge potential $\xi \in C^1(\Gamma; O(d))$ is synchronizable if and only if there exists $g \in C^0(\Gamma; O(d))$ such that $\xi_{ij} = g_i g_j^{-1}$ for all $(i, j) \in E$. The frustration $\nu(\Gamma)$ defined in (49) can thus be rewritten as

$$\nu(\Gamma) = \frac{1}{2d} \frac{1}{\text{vol}(\Gamma)} \inf_{g \in C^0(\Gamma; O(d))} \sum_{i,j \in V} w_{ij} \|g_i g_j^{-1} - \rho_{ij}\|_F^2 = \frac{1}{2d} \frac{1}{\text{vol}(\Gamma)} \inf_{\xi \in C_{\text{sync}}^1(\Gamma; O(d))} \sum_{i,j \in V} w_{ij} \|\xi_{ij} - \rho_{ij}\|_F^2,$$

where we define

$$C_{\text{sync}}^1(\Gamma; O(d)) := \{\xi \in C^1(\Gamma; O(d)) \mid \xi \text{ synchronizable}\} = \{\xi \in C^1(\Gamma; O(d)) \mid \text{Hol}_\xi(\Gamma) \text{ is trivial}\}.$$

Therefore, in the fibre bundle framework, a synchronization problem asks for a synchronizable edge potential that is “as close as possible” to a prescribed edge potential, or geometrically speaking, for a trivial flat bundle “as close as possible” to a given flat bundle. One approach, from the point of view of Proposition 2.1, is to find d linearly independent cochains in $C^0(\Gamma; \mathbb{R}^d)$ that are “as close as possible” to being a global frame of $\mathcal{B}_\rho[\mathbb{R}^d]$ in the sense of minimizing the *frustration of a \mathbb{S}^{d-1} -valued cochain*

$$\eta(f) = \frac{\langle f, \Delta_\rho^{(0)} f \rangle}{\|f\|^2} = \frac{\frac{1}{2} \sum_{i,j \in V} w_{ij} \|f_i - \rho_{ij} f_j\|^2}{\sum_{i \in V} d_i \|f_i\|^2} = \frac{1}{2\text{vol}(\Gamma)} [f]^\top L_1 [f], \quad \forall f \in C^0(\Gamma; \mathbb{S}^{d-1}), \quad \|f\| \neq 0,$$

which equals zero if and only if f defines a global section on $\mathcal{B}_\rho[\mathbb{R}^d]$ (the constraint $\|f\| \neq 0$ is also indispensable from a geometric point of view, as any vector bundle trivially admits the constant zero global section). This provides a geometric interpretation of the spectral relaxation algorithm in [BSS13]. From a perturbation point of view, the magnitudes of the smallest n eigenvalues of $\Delta_\rho^{(0)}$ measure the deviation from degeneracy of the d -dimensional eigenspace of lowest frequencies, and can thus be interpreted as the extent to which $\mathcal{B}_\rho[\mathbb{R}^d]$ deviates from admitting d linearly independent global sections and being a trivial bundle. The Cheeger-type inequality established in [BSS13] quantitatively confirms this geometric intuition relating $\nu(\Gamma)$ to the magnitude of d smallest eigenvalues of $D_1^{-1} L_1$ (the random walk version of the graph connection Laplacian):

$$\frac{1}{d} \sum_{k=1}^d \lambda_k(D_1^{-1} L_1) \leq \nu(\Gamma) \leq \frac{Cd^3}{\lambda_2(L_0)} \sum_{k=1}^d \lambda_k(D_1^{-1} L_1), \quad (50)$$

where $C > 0$ is a constant, $\lambda_2(L_0)$ is the spectral gap of Γ associated with the graph Laplacian L_0 , and $\lambda_k(D_1^{-1} L_1)$ is the k th smallest eigenvalue of $D_1^{-1} L_1$. (The actual version stated in [BSS13] is for the smallest d eigenvalues of the normalized graph connection Laplacian $D_1^{-1/2} L_1 D_1^{-1/2}$, but note that $D_1^{-1/2} L_1 D_1^{-1/2}$ has the same eigenvalues as $D_1^{-1} L_1$.)

Classical Cheeger inequalities [Che70, AKPW95, Chu07] relate isoperimetric constants or cuts on graphs and manifolds to the spectral gap of a graph Laplacian or Laplace-Beltrami operator. There have been Cheeger-type inequalities for simplicial complexes with the objective of understanding high-dimensional generalization of expander graphs [MS16, PR16, PRT15, SKM14]. These results are all concerned with partitioning graphs, manifolds, or simplicial complexes. The Cheeger-type inequality in equation (50) differs from standard Cheeger inequalities in that the cochains are group- or vector-valued. In addition, the frustration $\nu(\Gamma)$ is not related with any optimality for graph partitioning — in the sense of Proposition 2.1, $\nu(\Gamma)$ measures the triviality of a fibre bundle as a whole. The algorithm we will propose in Section 3 is an attempt to address the graph cut problem based on the synchronizability of the partitions resulted.

3 Learning Group Actions by Synchronization

In this section we specify an algorithm for learning group actions from observations based on synchronization. We also use simulations to provide some insight towards the performance of the algorithm.

3.1 Motivation and General Formulation

We first state some basic terminology from the general theory of group actions that will be used extensively. If G is a group and X is a set, a *left group action of G on X* is a map $\phi : G \times X \rightarrow X : (g, x) \mapsto \phi(g, x)$ such that

$$\phi(e, x) = x, \forall x \in X \quad \text{if } e \text{ is the identity element of } G$$

and

$$\phi(g, \phi(h, x)) = \phi(gh, x), \quad \forall x \in X, \forall g, h \in G.$$

To simplify notation, we will abbreviate $\phi(g, x)$ as $g.x$. The *orbit* of any element $x \in X$ under the action of G is defined as the set $G.x := \{g.x \mid g \in G\}$. If we introduce an equivalence relation on X by setting

$$x \sim y \iff x = g.y \text{ for some } g \in G,$$

then clearly $x \sim y$ if and only if $G.x = G.y$. The set X is naturally partitioned into the disjoint unions of orbits, and each orbit Y is an *invariant subset* of X under the action of G in the sense that $G.Y \subset Y$. If for any pair of distinct elements x, y of X there exists $g \in G$ such that $g.x = y$, we say that the action of G on X is *transitive*. Note that the total space X is an invariant subset in its own right, and the action of G on each orbit is obviously transitive. If the set X is finite and there exists a constant time procedure to verify whether any two elements are equivalent under transformations, the problem of partitioning X into disjoint subsets of orbits can be solved in polynomial time complexity with respect to the size of X .

In practice we are often interested in classification or clustering tasks which can be framed as follows: given a dataset $X = \{x_1, \dots, x_n\}$ of n objects, find a correspondence or transformation between each pair of distinct objects. We will see these pairwise correspondences often play the role of nuisance variables and one needs to “quotient out” the influence of these variables in downstream analysis (e.g. for most practical applications of synchronization problems [SZSH11, TSR11, BCSZ14] and alignment problems in statistical shape analysis [BPG⁺15]). The intuition as to why some of these pairwise correspondences are nuisance variables one can often with greater fidelity transform one object into another via intermediary transformations to other objects rather than a direct transformation between objects. Sometimes, for instance in the analysis of a collection of shapes in computer graphics [HZG⁺12, NBCW⁺11b, HG13, CGH14, KKBL15, MDK⁺16] and group-wise registration in automated geometric morphometrics [BLS⁺11, AADL13, LD11, LPD13, KH15], the pairwise transformations contains crucial information and are important on their own right. A common challenge in both of the above problems is that the fidelity of pairwise comparisons can be extremely variable over the data. We illustrate this challenge using the example of computing continuous Procrustes distances between disk-type shapes [AADL13] in automated geometric morphometrics. The core of the algorithm is an efficient strategy for searching the Möbius transformation group of the unit disk to obtain a diffeomorphism between the shapes that minimizes an energy functional. It has been observed that for similar shapes (in the sense of having a small pairwise distance), the resulting diffeomorphism is often of high quality and can reflect the correspondence of biological traits. If the shape pair is highly dissimilar, the diffeomorphism tends to suffer from various structural errors (see e.g. [GYDB] and [Gao15, Chapter 5]). Similar issues have also been observed in the field of non-rigid shape registration in geometric processing — successful feature extraction and matching techniques for near-isometric shapes abound [SOG09, ASC11, BK10, RBBK10, KBLB12, LZ14], whereas registering shape pairs with large deformation is still considered a difficult open problem [BBK08, KABL14, APL15]. Recently, a series of works [NBCW⁺11b, HZG⁺12, HG13, CGH14, KKBL15, MDK⁺16] proposed to jointly compute all pairwise correspondences within a collection subject to “consistency constraints” that require the composition of resulted maps along any cycle within the collection be approximately the identity map. The idea in this approach is that pairwise correspondences between dissimilar shapes are implicitly approximated by concatenating many correspondences between similar shapes with the individual correspondences have high fidelity, thus avoiding directly solving non-convex optimization problems with large numbers of local minimizers. Similar ideas can also be found in recent progress in automated geometric morphometrics where a *Minimum Spanning Tree* (MST) provides the concatenating of correspondences [BPG⁺15, GYDB]. It has been observed by morphologists that cycle-consistent constraints are more

often satisfied for a collection of samples within a species versus samples across a variety of species, suggesting that inconsistency may be used for species clustering.

Motivated by the above algorithms and approaches, we propose to study the following general problem of *Learning Group Actions* (LGA):

Problem 1 (Learning Group Actions). Given a group G acting on a set X , simultaneously learn a new action of G on X and a partition of X into disjoint subsets X_1, \dots, X_K , such that the new action is as close as possible to the given action and cycle-consistent on each X_i ($1 \leq i \leq K$).

The LGA problem can also be understood as a variant of the classical clustering problem, in which the coarse-graining is based on the cycle-consistency of group actions rather than pairwise similarity or spatial configuration of elements in the dataset. A solution of the LGA problem provides not only a partition of the input dataset but also cycle-consistent group actions within each cluster. It is useful to notice that all group elements implemented as pairwise actions within the same partition X_i form a subgroup of G ; the LGA problem can thus also be considered as “learning” subgroups of a prescribed “ambient group” that optimally fit a given dataset X . In other words, by solving an LGA problem we identify the “correct” transformation group for a dataset, which in most practical situations are much more tightly adapted to the given data than the potentially massive group of all possible transformations G .

Example 3.1. If the set X is a vector space and we seek a direct sum decomposition $X = \bigoplus_{i=1}^K X_i$ instead of a partition $X = \bigcup_{i=1}^K X_i$, the LGA problem reduces to the search for all irreducible G -subrepresentations of X .

Example 3.2. Consider a point set $X = \{x_1, \dots, x_n\}$ equipped with a labeling map $S : X \rightarrow \{\pm 1\}$ that assigns to each x_i either value $+1$ or -1 . We say x_i has *positive spin* if $S(x_i) = 1$ and has *negative spin* if $S(x_i) = -1$. Let $G = \{\pm 1\}$ act on X transitively as $(g_{ji}, x_i) \mapsto x_j$, $g_{ji} = S(x_j)S(x_i)$. Suppose the spin of each point in X (i.e. the label map S) is unknown, but we have full access to the group actions $\{g_{ij}\}$, we can reconstruct S — up to flipping labels ± 1 — by spectral clustering the dataset X , viewed as vertices of a complete graph Γ with weight $w_{ij} = g_{ij}$ on the edge connecting x_i and x_j . Under circumstances where some group actions g_{ji} are subject to a sign-flip error (noisy measurements), or/and the graph Γ is not complete (incomplete measurements), spectral or semi-definite programming relaxation techniques can still be used to recover S up to permuting labels ± 1 (see e.g. [CKS15]). With X and (potentially noisy and incomplete) $\{g_{ji}\}$ as input, this spectral clustering example can be considered as an instance of LGA: the output consists of a partition of X into positive/negative spin subsets, as well as the trivial subgroup $\{+1\}$ of $G = \{\pm 1\}$ acting in a cycle-consistent manner on both partitions.

Example 3.2 provides further motivation to consider a version of LGA in the context of synchronization problems. We are given a graph $\Gamma = (V, E)$ and the data X , where the vertex set V is identified with observations in X and the edges in E representing pairwise relations between elements of X . It is natural to consider a *partition of the graph* Γ in this setup decomposition of Γ into connected subgraphs such that the vertices of the subgraphs form a partition of the set of vertices of Γ .

Problem 2 (Learning Group Actions by Synchronization). Let $\Gamma = (V, E)$ be an undirected weighted graph, G a topological group, and $\rho \in C^1(\Gamma; G)$ a given edge potential on Γ . Furthermore, assume the vertex set V is equipped with a cost function $\text{Cost}_G : G \times G \rightarrow [0, \infty)$. Denote \mathcal{X}_K for all partitions of Γ into K nonempty connected subgroups ($K \leq n$) and

$$\nu(S_i) = \inf_{f \in C^0(\Gamma; G)} \sum_{j,k \in S_i} w_{jk} \text{Cost}_G(f_j, \rho_{jk} f_k), \quad \text{vol}(S_i) = \sum_{j \in S_i} d_j, \quad 1 \leq i \leq K.$$

Solve the optimization problem

$$\min_{\{S_1, \dots, S_K\} \in \mathcal{X}_K} \frac{\max_{1 \leq i \leq K} \nu(S_i)}{\min_{1 \leq i \leq K} \text{vol}(S_i)} \quad (51)$$

and output an optimal partition $\{S_1, \dots, S_K\}$ together with the minimizing $f \in C^0(\Gamma; G)$.

In the following, we shall refer to Problem 2 as *Learning Group Actions by Synchronization* (LGAS). When $G = O(d)$ or $U(d)$ and Cost_G is the squared Frobenius norm on $d \times d$ matrices, $\nu(S_i)$ is clearly the frustration (49) of the subgraph of Γ with vertices in S_i , up to a multiplicative constant depending only on Γ and dimension d . The minimizing vertex potential defines a synchronizable edge potential on the entire graph Γ , thus also gives rise to a cycle-consistent action on each partition. Note that the objective function (51) does not account for the discrepancy between the realized synchronizable edge potential and the original ρ on edges across partitions — intuitively, solving Problem 2 amounts to forming partitions by economically “dropping out” appropriate edges in Γ to minimize the total frustration.

3.2 SynCut: A Heuristic Algorithm for Learning Group Actions by Synchronization

In this subsection, we will investigate Problem 2 (LGAS) in the context of $O(d)$ -synchronization problems, focusing on the simpler setting where $K = 2$. In this case, (51) simplifies to

$$\min_{S \subset V} \frac{\max\{\nu(S), \nu(S^c)\}}{\min\{\text{vol}(S), \text{vol}(S^c)\}}.$$

Note that

$$\max\{\nu(S), \nu(S^c)\} \leq \nu(S) + \nu(S^c) \leq 2 \max\{\nu(S), \nu(S^c)\},$$

we can thus consider — drawing an analogy with the standard approach of studying Cheeger numbers through normalized cuts — the following optimization problem closely related with (51):

$$\begin{aligned} \xi_\Gamma &:= \min_{S \subset V} \xi(S) \\ &:= \min_{S \subset V} [\nu(S) + \nu(S^c)] \left(\frac{1}{\text{vol}(S)} + \frac{1}{\text{vol}(S^c)} \right). \end{aligned} \quad (52)$$

Recall from (49) that ξ_Γ further simplifies into

$$\begin{aligned} \xi_\Gamma &= \min_{S \subset V} \inf_{g \in C^0(\Gamma; O(d))} \frac{1}{2d} \frac{1}{\text{vol}(\Gamma)} \sum_{\substack{i, j \in V \\ (i, j) \notin \partial S}} w_{ij} \|g_i - \rho_{ij} g_j\|_F^2 \cdot \frac{\text{vol}(\Gamma)}{\text{vol}(S) \text{vol}(S^c)} \\ &= \min_{S \subset V} \inf_{g \in C^0(\Gamma; O(d))} \frac{1}{2d} \frac{1}{\text{vol}(S) \text{vol}(S^c)} \sum_{\substack{i, j \in V \\ (i, j) \notin \partial S}} w_{ij} \|g_i - \rho_{ij} g_j\|_F^2, \end{aligned} \quad (53)$$

where

$$\partial S := \{(u, v) \in E \mid u \in S, v \in S^c \text{ or } u \in S^c, v \in S\}.$$

In other words, the goal of solving the optimization problem (53) is to lower the total frustration of the graph Γ by dropping out a minimum set of edges under the constraint that the residual graph consists of two connected components; this is equivalent to say that we seek a most economic graph cut in terms of reducing total frustration. To simplify statements, we shall refer to $\|g_i - \rho_{ij} g_j\|_F^2$ as the *frustration on edge* $(i, j) \in E$ of vertex potential g with respect to the edge potential ρ , and call the collection of frustrations on all edges the *edge-wise frustrations*. The sum of all edge-wise frustrations will be referred to as the *total frustration*.

Formulation (53) motivates a greedy algorithm that alternates between minimizing graph cuts and vertex potentials. We shall refer to this algorithm as *Synchronization Cut*, or SynCut for short; see Algorithm 1. We describe the main steps in SynCut below:

Step 1. Initialization: Input data include the weighted graph $\Gamma = (V, E, w)$, edge potential $\rho \in C^1(\Gamma; G)$, and parameters required for the spectral clustering subroutine plus termination conditions for the main loop. Initialize iteration counter $t = 0$, and dynamic graph weights ϵ to be the input graph weights w ;

Step 2. Global Synchronization: Synchronize the edge potential ρ on the entire graph Γ with respect to edge weights. Any synchronization algorithm can be used in this step, e.g. spectral relaxation [BSS13, CKS15] or SDP relaxation [BCSZ14, Sin11, CKS15, BKS16, NRV13]. Note that in this step the synchronization is performed on Γ with dynamic weights ϵ instead of the original weights w . Denote $f^{(t)}$ as the edge potential on Γ at the edge potential at the t -th iteration;

Step 3. Spectral Clustering (First Pass): Update dynamic weights ϵ based on the frustration of $f^{(\ell)}$ on each edge by

$$\epsilon_{ij} = w_{ij} \exp\left(-\frac{1}{\sigma} \left\| f_i^{(\ell)} - \rho_{ij} f_j^{(\ell)} \right\|_F^2\right), \quad \text{where } \sigma > 0 \text{ is the average of all non-zero edge-wise frustrations}$$

then partition the vertex set V of graph Γ into K clusters S_1, \dots, S_K using spectral clustering based on the updated dynamic weights ϵ . The goal is to cut the graph Γ into more synchronizable clusters; edges causing large frustration are assigned relatively smaller weights ϵ_{ij} to increase the chance of being cut. To simplify notation, we will also use S_ℓ ($1 \leq \ell \leq K$) to denote the subgraph of Γ spanned by the vertices in S_ℓ ;

Step 4. Local Synchronization: Synchronize the edge potential ρ within each partition S_ℓ , $1 \leq \ell \leq K$. If we denote $\rho|_{S_\ell}$, $\epsilon|_{S_\ell}$ for the restrictions of ρ , ϵ to S_ℓ , respectively, then this step solves the synchronization problem on weighted graph $(S_\ell, \epsilon|_{S_\ell})$ for prescribed edge potential $\rho|_{S_\ell}$. Again, any synchronization algorithm can be used in this step. Denote $g^{(\ell)}$ for the resulting vertex potential on S_ℓ ;

Step 5. Collage: After obtaining $g^{(\ell)}$ for each local synchronization on S_ℓ , we make a ‘‘collage’’ from these local solutions to form a global vertex potential defined on the entire graph Γ . Since each $g^{(\ell)}$ is obtained from synchronizing within S_ℓ , the collected local solutions $\{g^{(\ell)}\}_{\ell=1}^K$ generally incur large incompatibility (frustration) on edges across partitions. Our strategy is to find K elements $h_1, \dots, h_K \in G$, where each h_ℓ acts on $g^{(\ell)}$ by $g_u^{(\ell)} \mapsto g_u^{(\ell)} h_\ell$, $\forall u \in S_\ell$, such that the *total cross-partition frustration*

$$C(h_1, \dots, h_K \mid \{S_\ell\}_{1 \leq \ell \leq K}, \{g^{(\ell)}\}_{1 \leq \ell \leq K}) := \sum_{1 \leq p \neq q \leq K} \sum_{\substack{(u,v) \in E \\ u \in S_p, v \in S_q}} w_{uv} \left\| g_u^{(p)} h_p - \rho_{uv} g_v^{(q)} h_q \right\|_F^2 \quad (54)$$

is minimized. Note that this is essentially synchronizing an edge potential

$$\tilde{\rho}_{pq} = \begin{cases} \sum_{\substack{(u,v) \in E \\ u \in S_p, v \in S_q}} w_{uv} (g_u^{(p)})^{-1} \rho_{uv} g_v^{(q)} & \text{if partitions } S_p, S_q \text{ are connected} \\ 0 & \text{otherwise} \end{cases} \quad (55)$$

on a reduced complete graph $\tilde{\Gamma}_K$ consisting of K vertices where each vertex represent one of the K partitions S_1, \dots, S_K . It thus simply requires calling the synchronization routine again to obtain h_1, \dots, h_K , but this time the scale of the synchronization problem is often much smaller than the previous global and local synchronization steps. Also note that for the binary cut case $K = 2$ and $G = O(d)$ this collage step is even simpler: it suffices to perform an single SVD on the $d \times d$ matrix

$$\sum_{\substack{(u,v) \in E \\ u \in S_1, v \in S_2}} w_{uv} (g_u^{(p)})^{-1} \rho_{uv} g_v^{(q)} = U \Sigma V^\top$$

and set $h_1 = UV^\top$, $h_2 = I_{d \times d}$.

Step 6. Spectral Clustering (Second Pass): Update dynamic weights ϵ based on the frustration of f^* on each edge by

$$\epsilon_{ij} = w_{ij} \exp\left(-\frac{1}{\sigma} \left\| f_i^* - \rho_{ij} f_j^* \right\|_F^2\right), \quad \text{where } \sigma > 0 \text{ is the average of all non-zero edge-wise frustrations}$$

then partition Γ into K clusters S_1, \dots, S_K using spectral clustering for a second time, based on the updated dynamic weights ϵ .

Step 7. Repeat Step 2 to Step 7 Until Convergence. The termination condition can be specified either by a maximum number of iterations or monitoring the change of the quantity

$$\xi(\{S_1, \dots, S_K\}) := \left(\sum_{\ell=1}^K \nu(S_\ell) \right) \left(\sum_{k=1}^K \frac{1}{\text{vol}(S_k)} \right). \quad (56)$$

At the end of the procedure, return the partitions $\{S_1, \dots, S_K\}$ and the final edge potential f^* from the most recent updates. The cycle-consistent edge potential on partition S_ℓ is encoded in the restriction of f^* to S_ℓ .

Algorithm 1 SYNCHRONIZATION CUT: Learning Group Actions by Synchronization

```
1: procedure SYNcut( $\Gamma, \rho, K$ ) ▷ weighted graph  $\Gamma = (V, E, w), \rho \in C^1(\Gamma; G)$ , number of partitions  $K$ 
2:    $t = 0$ 
3:    $\epsilon = w$ 
4:   while not converge do
5:      $f^{(t)} \in C^0(\Gamma; G) \leftarrow \text{SYNCHRONIZE}(\Gamma, \rho, \epsilon)$ 
6:      $\sigma \leftarrow$  average non-zero edge-wise frustrations of  $f^{(t)}$ 
7:     for  $(i, j) \in E$  do ▷ calculate weights  $\epsilon$  on graph  $\Gamma$  for spectral clustering
8:        $\epsilon_{ij} \leftarrow w_{ij} \exp\left(-\frac{1}{\sigma} \|f_i^{(t)} - \rho_{ij} f_j^{(t)}\|_F^2\right)$ 
9:     end for
10:     $\{S_1, \dots, S_K\} \leftarrow \text{SPECTRALCLUSTERING}(\Gamma, \epsilon)$ 
11:    for  $\ell = 1, \dots, K$  do
12:       $g^{(\ell)} \in \Omega^0(S_\ell; G) \leftarrow \text{SYNCHRONIZE}(S_\ell, \rho|_{S_\ell}, \epsilon|_{S_\ell})$ 
13:    end for
14:     $f^* \in \Omega^0(\Gamma; G) \leftarrow \text{COLLAGE}\left(\{S_\ell\}_{\ell=1}^K, \{g^{(\ell)}\}_{\ell=1}^K\right)$ 
15:     $\sigma \leftarrow$  average non-zero edge-wise frustrations of  $f^*$ 
16:    for  $(i, j) \in E$  do ▷ update weights  $w^{(t)}$  on graph  $\Gamma$  for next iteration
17:       $\epsilon_{ij} \leftarrow w_{ij} \exp\left(-\frac{1}{\sigma} \|f_i^* - \rho_{ij} f_j^*\|_F^2\right)$ 
18:    end for
19:     $\{S_1, \dots, S_K\} \leftarrow \text{SPECTRALCLUSTERING}(\Gamma, \epsilon)$ 
20:     $t \leftarrow t + 1$ 
21:  end while
22:  return  $\{S_1, \dots, S_K\}, f^*$  ▷  $f^*$  defines a cycle-consistent edge potential on each partition
23: end procedure
```

3.3 Results on Simulated Random Synchronization Networks

In this subsection, we use simulations to provide some intuition for the behavior of SynCut under the setting $K = 2$ (two partitions). We first specify a random procedure to simulate input data — a connected random graph with a prescribed edge potential — for synchronization problems. In addition, the random graph generation procedure will be controlled by a parameter that allows us to adjust the level of obstruction to the synchronizability of the prescribed edge potential over the generated graph. We then specify the metrics used for performance measure. We conclude by demonstrating that the partition generated from SynCut recovers the two synchronizable connected components with high accuracy and within relatively few numbers of iterations. For the simplicity of statements, we refer to each pair of generated graph and edge potential an instance of a *random synchronization network*.

3.3.1 Random Synchronization Network Simulation

We first specify the procedure to generate the random graphs. Our intention is to sample random graphs with sufficiently variable spectral gaps, based on the intuition that a large spectral gap of the underlying graph results in greater obstruction to the synchronizability of the edge potential constructed by the procedures that will soon be described in this subsection. We first generate two partitions S_1, S_2 with an equal number of vertices. Each partition is a connected component built from a vertex degree sequence of random integers uniformly distributed in an interval (say 5 to 8), adapting an algorithm first proposed in [BD10]; when the interval is a single integer, the connected component is a regular graph. Random edges are then added to link the two partitions S_1, S_2 . The number of inter-component random edges positively correlates with the spectral gap, as shown in Figure 1, suggesting that this number can be used as a parameter to adjust the level of obstruction to cutting the graph into two connected components S_1 and S_2 .

A subtlety in this random network generation procedure is that a uniform distribution on the number of inter-component links does not induce a uniform distribution on the spectral gaps of the generated random graphs, due to concentration effects. A precise characterization of the distribution of spectral gaps in our random graph model is interesting on its own right but beyond the scope of this paper. We refer interested readers to the existing literature on

the spectral gaps of random graphs such as [CLV03, COL06, HKP12]. In practice, we simply use a large number of random trials to generate sufficiently many sample graphs with spectral gaps within desirable ranges; see Figure 2a.

After drawing an instance of the random graph, we randomly construct an edge potential that is synchronizable within S_1 and S_2 , but not necessarily synchronizable on the inter-component links. The procedure to generate the random edge potentials proceeds as follows:

- (1) Randomly generate a vertex potential $g \in C^0(\Gamma; G)$ for the entire graph Γ ;
- (2) Set the value of ρ on edge (i, j) according to

$$\rho_{ij} = \begin{cases} g_i g_j^{-1} & \text{if both } i, j \in S_1 \text{ or } i, j \in S_2, \\ \text{a random matrix in } O(d) & \text{otherwise.} \end{cases}$$

The vertex potential $g \in C^0(\Gamma; G)$ will no longer attain the minimum frustration for the entire graph with respect to the prescribed edge potential ρ , due to the edges added between the partitions that are much less likely synchronizable.




Figure 1: A scatter plot displaying the correlation between the number of inter-component links and the spectral gap in our random graph model, with $N = 100$ vertices and the (integer) number of inter-component links uniformly distributed between 100 and 250.

We consider each run of SynCut as successful if both output partitions are synchronizable connected components, i.e. if SynCut recovers the original partitions S_1, S_2 . The performance of SynCut is measured using the *error ratio* computed by dividing the number of erroneously clustered vertices by the total number of vertices. If SynCut successfully recovers S_1, S_2 , this error ratio is 0; if the output partition is close to a random guess, or if the algorithm fails to separate the vertices into distinct clusters, the error ratio is 0.5. The error ratios of the partitions output from SynCut are then compared with a baseline graph cutting algorithm using normalized graph cut (NCut) that does not utilize any information in the prescribed edge potential; see e.g. [SM00, Von07].

3.3.2 Simulation Results

In this simulation study, we set the number of vertices in each of the two synchronizable components to $N = 100$, and the entries of the vertex degree sequence of random integers are independently uniformly distributed between 4 and 8. The number of inter-component links between the two synchronizable components is drawn uniformly between 100 and 250. The edge potentials are valued in the orthogonal group $O(d)$ with $d = 5$. We terminate SynCut either after 10 iterations or if the change in the value of the objective function ξ [see (56)] between consecutive iterations falls below a preset tolerance of 10^{-8} . We plot in Figure 2(a) the spectral gaps of 10,000 realizations of our random network model. In Figure 2(b) and Figure 2(c) we observed that the error ratios in these 10,000 runs of SynCut tend to be much smaller than NCut, suggesting that SynCut outputs more accurate partitions with respect to synchronizability. In Figure 2(e), we again see that SynCut outperforms NCut and the amount of improvement increases with the magnitude of the spectral gap. Figure 2(d) shows that SynCut converges quickly.

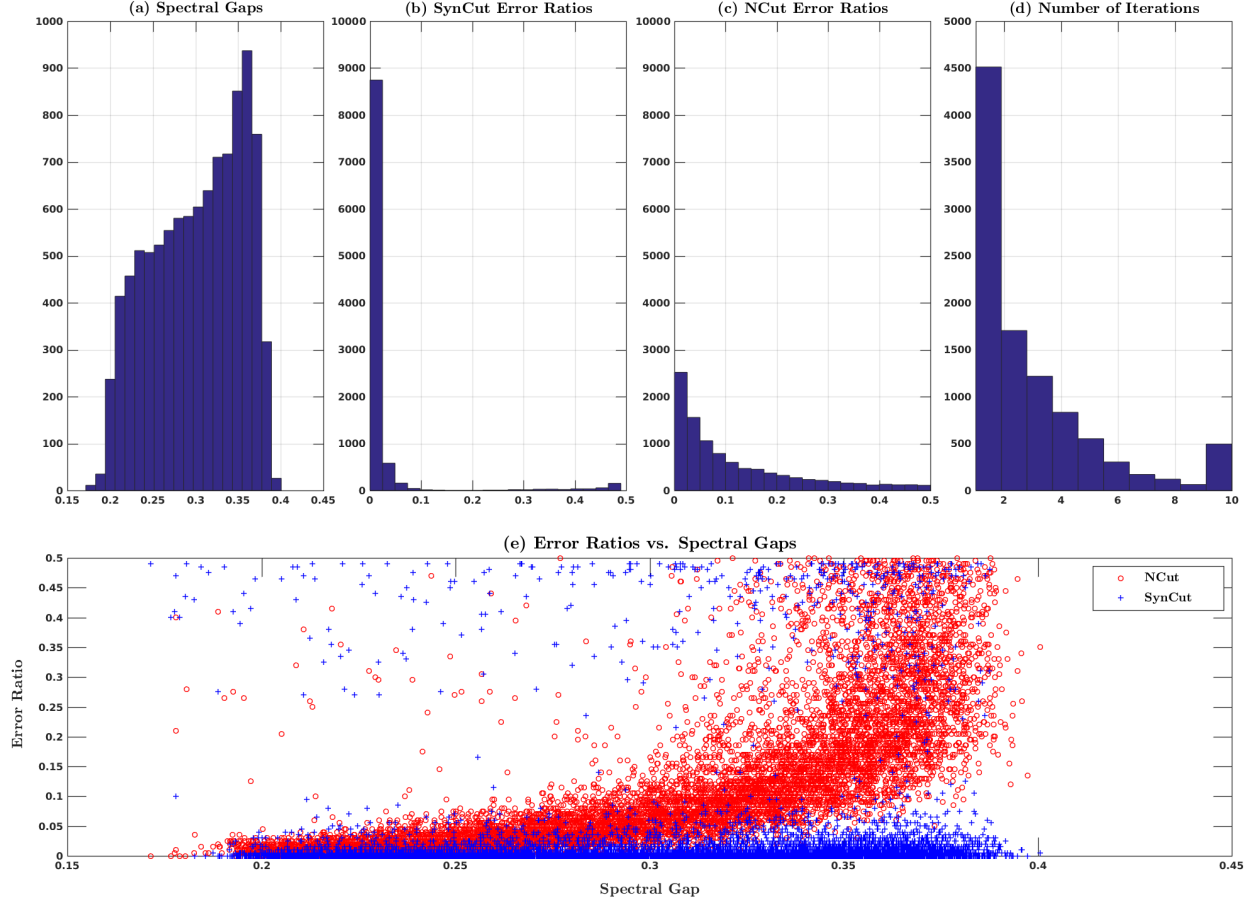


Figure 2: (a) Histogram of the spectral gap of the 10,000 random graphs drawn from our model. (b) Histogram of the error ratios of the SynCut clustering results. (c) Histogram of the error ratios of the baseline NCut clustering results. (d) Histogram of the number of iterations for SynCut. (e) Scatter plots of the error ratios of SynCut and NCut versus spectral gap.

We now focus on a particular instance of a random synchronization network to better understand SynCut in comparison with the spectral relaxation algorithm proposed in [BSS13]. Each synchronizable component in the random synchronization network shown in Figure 3(a) is a regular graph containing $N = 100$ vertices and 250 edges, generated with a constant vertex degree sequence of 5. We color the edges within and between synchronizable components in blue and red, respectively. In Figure 3(b) we plot the edge-wise frustration for the vertex potential g used to generate the edge potential ρ prescribed to the network. As expected, the frustration is zero within each synchronizable component but large on the edges across components. Figure 3(c) and Figure 3(d) show the edge-wise frustrations for two vertex potentials obtained from the spectral relaxation algorithm [BSS13] and SynCut, respectively. Though the total frustration is larger for SynCut than spectral relaxation, the SynCut solution concentrates most of the frustration on the non-synchronizable inter-component edges, with a distribution of edge-wise frustrations closer to the distribution for the initial vertex potential g . This suggests that applying a spectral graph cut algorithm using the edge-wise frustration of the SynCut solution as a dissimilarity measure is advantageous, as the distribution in Figure 3(d) identifies the obstructions to synchronizability in the synchronization network more accurately.




Figure 3: (a) A random graph consisting of two synchronizable connected components, each with 100 vertices and 250 edges (in blue), and 100 non-synchronizable inter-component edges (in red). The edge potential takes value in the orthogonal group $O(5)$. All vertex degrees in each synchronizable component are set to 5. (b) Edge-wise frustration for the vertex potential g used to generate the prescribed edge potential ρ . As expected, frustration is small within each connected component but large between components. (c) Edge-wise frustration for the vertex potential obtained from spectral relaxation [BSS13]. The total frustration is lower than that in the top right figure, but the inter-component edges carries relatively lower frustration since the relaxation procedure tends to “spread” the non-synchronizability across the entire graph. (d) Edge-wise frustration for the vertex potential obtained from SynCut. The total frustration is higher than that for the spectral relaxation solution, but the distribution of frustrations on the edges is closer to that of vertex g and can thus be used to recover the synchronizable connected components.

4 Application to Automated Geometric Morphometrics

In this section we formulate a problem in automated geometric morphometrics in terms of LGAS, then apply the SynCut algorithm to provide a solution. In Section 4.1 we provide some background in geometric morphometrics and its relation to synchronization problems. In Section 4.2 we apply SynCut to a collection of second mandibular molars of prosimian primates and non-primate close relatives. The morphological hypothesis is that the geometric

traits of second mandibular molars cluster into 3 dietary regimens: *folivorous* (herbivores that eat leaves), *frugivorous* (herbivores or omnivores that prefer fruit), and *insectivorous* (a carnivore that eats insects). We will show the SynCut result, which is based on the synchronizability of pairwise correspondences, and compare it with a distance-based clustering result using diffusion maps [CL06].

4.1 Geometric Morphometrics and Synchronization

The classic tool in geometric morphometrics is *Procrustes analysis*. The basic assumption underlying this analysis framework is that most of the geometric information on each shape can be efficiently encoded in a set of landmark points carefully picked to highlight the morphometrical phenotypes (variation in the geometric shape of an organism). The *Procrustes distance* between two shapes is the average Euclidean distance between corresponding landmarks, after applying a rigid motion (rotations, reflections, translations, and their compositions) to optimally align the two sets of landmarks. If all the shapes are marked with an equal number of landmarks but the landmark correspondence is not known *a priori*, a combinatorial search can be performed over all permutations of one-to-one landmark correspondences, and the minimum average Euclidean distance between corresponding landmarks can be taken as a dissimilar measure between the two shapes. Comparing a pair of shapes in this framework thus yields abundant pairwise information, including a scalar dissimilarity score, a rigid motion, and a permutation matrix encoding the one-to-one landmark correspondence.

In automated geometric morphometrics, landmark points are not used to represent the shapes, and algorithms search for an “optimal” transformation between a pair of whole shapes directly by minimizing energy functionals over a set of admissible transformations. Depending on the specific class of transformations and energy functional, the pairwise comparisons produce different types of correspondences between surfaces, such as conformal/quasiconformal transformations, isometries, area-preserving diffeomorphisms, or even transport-plans between surface area measures in a Wasserstein framework. Regardless of the type of admissible transformations, the algorithm can output a rigid motion for the optimal alignment between two shapes, as well as a dissimilarity or similarity score for such an alignment. See Figure 4 for an example of representing a collection of shapes using landmarks versus triangular meshes.

When the analysis is extended from comparing a single pair to a large collection of shapes, a crucial premise for downstream statistical analysis (e.g. *General Procrustes Analysis* (GPA) [Gow75, DM98, GD04]) is that the pairwise correspondences be *cycle-consistent*, meaning that propagating any landmark on any shape by consecutive correspondences along a close cycle of shapes should land exactly at the original landmark. Traditional landmark-based Procrustes analysis begins with consistently picking an equal number of landmarks on each shape, resulting in a large amount of pairwise correspondence relations that are cycle-consistent by construction. This is, however, not the situation with automated geometric morphometrics, where the correspondence transformations produced by automated algorithms are rarely cycle-consistent, even when one localizes the transformations within relatively “stable” regions where landmarks are affixed with the knowledge of an experienced geometric morphometrician. The necessity of cycle-consistent correspondences links automated geometric morphometrics to synchronization problems. An automated geometric morphometric algorithm will output for each pair of shapes a triplet consisting of a dissimilarity score, a rigid motion, and a pairwise transformation. We can use the dissimilarity scores to define a weighted graph Γ that captures the similarities within the collection, both qualitatively and quantitatively, by adjusting the number of nearest neighbors of each vertex and the weights on each edge. The rigid motions and pairwise transformations define two edge potentials on Γ , taking values in different groups. We list below some interesting synchronization problems arising from this formulation:

Three-Dimensional Euclidean Group. The rigid motions R_{ij} between shapes S_i and S_j that share an edge in Γ define an edge potential $R \in C^1(\Gamma; E(3))$, where $E(3)$ is the three-dimensional Euclidean group. Solving an $E(3)$ -synchronization problem over Γ with respect to R results in a globally consistent alignment for a collection of shapes, which is often crucial for initializing geometric morphometrical analysis algorithms such as *Dirichlet Normal Energy* [BBL⁺11], *Orientation Patch Analysis* [EWFJ07], and *Relief Index* [Boy08]. Algorithms that automatically align a collection of anatomical shapes in a globally consistent manner can also be viewed as primitive approaches for solving $E(3)$ -synchronization problems; see e.g. [Pue13, BPG⁺15, PWM⁺16, GBAD⁺16].

Orthogonal Group and Orientation Detection. If the shapes are preprocessed to superimpose the centers of mass at the same point, the translation component of each R_{ij} output from a pairwise landmark-based Procrustes analysis vanishes³. This reduces the global alignment problem to standard synchronization problems over the compact Lie group $O(3)$. Spectral and semidefinite programming (SDP) relaxation methods can then be applied directly to

³Note this is not the case for jointly analyzing a collection of shapes in a landmark-based Procrustes analysis framework; see e.g. [CKS15].

solve the global alignment problem. If we consider the edge potential $\rho \in C^1(\Gamma; \mathbb{Z}_2)$ defined by $\rho_{ij} = \det R_{ij}$, a \mathbb{Z}_2 -synchronization solution can be used to either partition the dataset into “left-handed” and “right-handed” subsets or conclude that such an orientation-based partition does not exist. We stated a similar situations in Example 3.2; other examples in this setting can be found in applications of *Orientable Diffusion Maps* [SW11].

Automorphism Groups. Certain classes of transformations C_{ij} between each pair of shapes S_i, S_j give rise to an edge potential on the graph Γ valued in an automorphism group of a canonical domain. For instance, algorithms such as *Möbius Voting* and the *Continuous Procrustes Distance* [AADL13] between disk-type surfaces rely on the computation of conformal maps between two shapes, based on uniformization parametrization techniques [PP93, AHTK99] that map each surface conformally to a canonical unit disk on the plane. By intertwining C_{ij} with the parametrizations of the source and the target shape, the correspondence between S_i and S_j can be equivalently considered as an element of the conformal automorphism group $\text{Aut}(\mathcal{D})$ of the planar unit disk \mathcal{D} . The group $\text{Aut}(\mathcal{D})$ is isomorphic to the *projective special linear group* $\text{PSL}(2, \mathbb{R})$, a non-compact simple real Lie group that is equivalent to the quotient of the *special linear group* $\text{SL}(2, \mathbb{R})$ by $\{\pm I_2\}$, where I_2 denotes the 2×2 identity matrix. Synchronization problems over $\text{PSL}(2, \mathbb{R})$ or $\text{SL}(2, \mathbb{R})$ require non-trivial extensions of the non-unique games (NUG) framework [BCS15] over compact Lie groups.

Groupoids. Other types of transformations C_{ij} between each pair of shapes S_i, S_j require further generalizations of the synchronization framework to edge potentials taking values in a *groupoid* rather than in a group. As an example, consider surface registration techniques based on area-preserving maps [AADL13, ZHT03, ZSG⁺13, SZS⁺13]. These techniques use conformal or area-preserving parametrizations to push forward surface area measures on S_i, S_j to measures μ_i, μ_j on the planar unit disk \mathcal{D} , respectively, then solve for a *transport map* on \mathcal{D} that pulls back μ_j to μ_i (or equivalently μ_i to μ_j). To formulate such “transport-map-valued” edge potentials in a synchronization framework, an edge potential should be allowed to take values in different classes of maps on different edges, with the only constraint that maps on consecutive edges can be composed; these ingredients have much in common in spirit with *fundamental groupoids* [Bro67, BHS11] and Haefliger’s *complexes of groups* [Hae91, Hae92]. Such a generalized framework for synchronization problems can also be used to analyze correspondences $\{C_{ij}\}$ that are *soft maps* [RCB97, SNB⁺12] or *transport plans* [LD11, LPD13, LZ14], where one replaces the set of transport maps between μ_i and μ_j with (probabilistic) couplings $\Pi(\mu_i, \mu_j)$ as in Kantorovich’s relaxation to the Monge optimal transport problem [Vi03, Vi08]. The *Horizontal Diffusion Maps* (HDM) framework [Gao16] and the application in automated geometric morphometrics [Gao15] are among the initial attempts in this direction.

4.2 Clustering Lemurs by Dietary Regimens using Synchronizability of Molar Surfaces

We focus on a real anatomical surface mesh dataset of second mandibular molars from 5 genera of prosimian primates and nonprimate close relatives. There are a total of 50 molars with 10 specimens from each genus. The five genera divide into three dietary regimens: the *Alouatta* and *Brachyteles* are folivorous, the *Ateles* and *Callicebus* are frugivorous, and the *Saimiri* are insectivorous. In Figure 4 we display four lemur molars from an anatomical surface dataset first published in [BLS⁺11], together with landmarks on each molar placed by evolutionary anthropologists. Similar datasets have been studied in a series of papers developing algorithms for automatic geometric morphometrics [BLS⁺11, AADL13, LD11, LPD13, KH15]. The chewing surface of each molar is digitized as a two-dimensional triangular mesh in \mathbb{R}^3 of disk-type topology (i.e. conformally equivalent with a planar disk). We will apply SynCut to these 50 molars and examine if the clustering is consistent with dietary regimens.

Method We first pre-process the dataset by translating and scaling each shape so that all surface meshes center at the origin and enclose unit surface area. We then apply the continuous Procrustes distance algorithm for each pair of teeth, generating a distance score d_{ij} , a diffeomorphism C_{ij} , and an orthogonal matrix $R_{ij} \in O(3)$ that optimally aligns S_j to S_i with respect to the diffeomorphism C_{ij} . We use the distance scores to construct a weighted K -nearest-neighbor graph Γ . The weights are defined as $w_{ij} = \exp(-d_{ij}^2/\sigma^2)$ with the bandwidth parameter $\sigma > 0$ set to be of the order of the average smallest non-zero distances. We apply SynCut to the edge potential $\rho \in C^1(\Gamma; O(3))$ defined by the alignments R_{ij} on the weighted graph (Γ, w) . Finally, we compare the clustering performance of SynCut with applying diffusion maps and spectral clustering directly to the weighted graph without the alignment information.

Results SynCut and diffusion maps both require the choice of a parameter K determining the number of nearest neighbors in the construction of the graph Γ . When $6 \leq K \leq 10$ both procedures accurately cluster the 50 molars in




Figure 4: The meshes of four lemur molars from an anatomical surface dataset first published in [BLS⁺11]. The colored dots on the molars are landmark points where identical colors indicate corresponding landmarks.




Figure 5: Consistent alignment of 50 lemur teeth based on applying SynCut to all pairwise alignments from the continuous Procrustes analysis [AADL13]. Each row corresponds to teeth from a genus, from top to bottom: *Alouatta*, *Ateles*, *Brachyteles*, *Callicebus*, *Saimiri*.

the dataset into the three distinct dietary regimens, see Figure 6 for the two-dimensional embedding plots for $K = 7$. SynCut produces slightly tighter and more distinguishable species clusters. Not surprisingly, for $K > 10$ — when the number of nearest neighbors exceeds the number of specimens in each genus — both algorithms are less accurate as K increases, with the accuracy of SynCut dropping faster than diffusion maps. This empirical observation is consistent with our intuition that the performance of SynCut is more sensitive to increased spectral gaps than diffusion maps.




Figure 6: Embeddings of the 50 lemur teeth dataset into \mathbb{R}^2 , obtained by applying diffusion maps (left) and SynCut (right) to the 7-nearest-neighbor graph. Both plots are post-processed using t-SNE [MH08]. (a) Diffusion maps applied to the weighted graph (Γ, w) successfully distinguishes three diet groups, but the genera are less distinguishable. (b) SynCut produces an edge-wise frustration matrix after the final iteration that can be used by diffusion maps to generate a low-dimensional embedding, in which both dietary groups and genera are more distinguishable.

5 Conclusion and Discussion

We provided in this paper a geometric framework for synchronization problems. We first related the synchronizability of an edge potential on a connected graph to the triviality of a flat principal bundle over the topological space underlying the graph, then characterized synchronizability from two aspects: the holonomy of the principal bundle, and the twisted cohomology of an associated vector bundle. On the holonomy side, we established a correspondence between two seemingly distant objects on a connected graph Γ , namely, the orbit space of the action of G -valued vertex potentials on G -valued edge potentials, and the representation variety of the fundamental group of Γ into G ; on the cohomology side, we built a twisted de Rham cochain complex on an associated vector bundle $\mathcal{B}_\rho[F]$ of the synchronization principal bundle \mathcal{B}_ρ , of which the zero-th degree cohomology group characterizes the obstruction to the synchronizability of the prescribed edge potential.

With the presence of a metric on the associated vector bundle $\mathcal{B}_\rho[F]$, we also developed a twisted Hodge theory on graphs. Independent of the contribution to synchronization problems, this theory is both a discrete version of the Hodge theory of elliptic complexes and a fibre bundle analogue of the discrete Hodge theory on graphs. Specifically for synchronization problems, this twisted Hodge theory realizes the graph connection Laplacian operator as the zero-th degree Hodge Laplacian in the twisted de Rham cochain complex. A Hodge-type decomposition theorem is also proven, stating that the image of the twisted codifferential is the orthogonal complement of the linear space of F -valued synchronization solutions, with respect to the bundle metric.

Motivated by the geometric intuitions gained from these theoretical results, we coined the problem of learning group actions (LGA), and proposed a heuristic algorithm, which we referred to as SynCut, based on iteratively applying synchronization and spectral graph techniques. Numerical simulations on synthetic and real datasets indicated that SynCut has the potential to cluster a collection of objects according to the synchronizability of a subset of partially observed pairwise transformations.

We conclude this paper by listing several problems of interest for future exploration. These are only a subset of a vast collection of potential directions concerning the mathematical, statistical, and algorithmic aspects of synchronization problems:

- 1) *The Representation Variety of Synchronization Problems.* When a prescribed edge potential ρ is not synchronizable over graph Γ , the goal of the synchronization problem is to find a synchronizable edge potential $\tilde{\rho}$ that is as close as possible to ρ in a sense that has been made clear in this paper. The point of view adopted in Section 2.1 is that the synchronization problem essentially concerns the orbits of ρ and $\tilde{\rho}$ under the action of all vertex potentials. It is natural to conceive a synchronization algorithm based on the geometry of the orbit space $C^1(\Gamma; G)/C^0(\Gamma; G)$ that enables efficiently “moving across” the orbits. Since the fundamental group of any connected graph is simply a free product of copies of \mathbb{Z} , we expect the representation variety $\text{Hom}(\pi_1(\Gamma), G)/G$ to possess relatively simple structures that could be used for guiding the design of novel synchronization algorithms with provable guarantees.
- 2) *Higher-order Synchronization Problems.* As a simplicial complex, the graph Γ only has 0- and 1-simplices, which results in only one cohomology group of interest in the de Rham cochain complex (41). By extending the twisted de Rham and Hodge theory developed in Section 2.2 to simplicial complexes of higher dimensions, we expect higher-order synchronization problems can be formulated and studied using tools and insights from high-dimensional expanders and the Hodge theory of elliptic complexes. Generalizing the current regime of synchronization problems, in which only pairwise transformations are considered, the higher-order analogies would enable the study of relations and interactions among multiple vertices in the graph Γ , which potentially opens doors towards higher-order graphical models and related statistical inference questions as well.
- 3) *Hierarchical Partial Synchronization Algorithms with Provable Guarantees.* The SynCut algorithm we proposed in this paper can be understood as an iterative hierarchical partial synchronization algorithm, based on the assumption that edge-wise synchronization is an indicator of the synchronizability of a prescribed edge potential over a proper subgraph. The numerical experiments on synthetic and real datasets suggested the validity of this intuition under our random graph model, but no provable guarantees exist either for the convergence or the effectiveness of algorithms similar or related to SynCut, to the best of our knowledge. Building a Cheeger-type inequality as the performance guarantee for SynCut attracted our attention, but even the analogy of Cheeger number (or graph conductance) in the setting of SynCut or LGA is not clear — whereas the Cheeger number depends only on the graph weights, which are fixed numbers on each edge independent of the graph cut, the notion of edge-wise frustration is highly non-local as the frustration depends on the behavior of the synchronization solution on the entire graph. We conjecture that a Cheeger-type inequality for SynCut, if exists, will reflect the global geometry information encoded by geometric quantities associated with the fibre bundle.
- 4) *Statistical Framework for Learning Group Actions.* The LGA problem presented in this paper is not formulated with a natural generative model for the dataset of objects with pairwise transformations; nor is assumed any concrete noise models. It would be of interest to provide a systematic, statistical framework under which the problem of LGA and LGAS can be quantitatively analyzed and understood; we believe such a framework also has the potential to bridge statistical inference with synchronization problems.

A Proofs of Proposition 1.1 and Formula (37)

Proof of Proposition 1.1. The construction of \mathcal{U} using the stars of the vertices of Γ ensures that

- (1) $U_i \cap U_j \neq \emptyset$ if and only if $(i, j) \in E$;
- (2) $U_i \cap U_j \cap U_k \neq \emptyset$ if and only if the 2-simplex (i, j, k) is in X .

For such pair (i, j) , define constant map $g_{ij} : U_i \cap U_j \rightarrow G$ as

$$g_{ij}(x) = \rho_{ij} \quad \forall x \in U_i \cap U_j.$$

Set $g_{ii} = e$ for all $1 \leq i \leq |V|$, and note that $g_{ij}(x) = g_{ji}^{-1}(x)$ for all $x \in U_i \cap U_j$ by our assumption on ρ . If ρ is synchronizable over G , let $f : V \rightarrow G$ be a vertex potential satisfying ρ , then $\rho_{ij} = f_i f_j^{-1}$ for all $(i, j) \in E$ from (1).

Thus $\rho_{kj}\rho_{ji} = \rho_{ki}$ for any triangle (i, j, k) in Γ , or equivalently that $g_{kj}(x)g_{ji}(x) = g_{ki}(x)$ for all $x \in U_i \cap U_j \cap U_k$. Therefore, $\{g_{ij} \mid 1 \leq i, j \leq |V|\}$ defines a system of *coordinate transformations* [Ste51, §2] with values in G . These data determine a principal fibre bundle \mathcal{P}_ρ with base space X and structure group G — by a standard construction in the theory of fibre bundles (see e.g. [Ste51, §3.2]) — of which local trivializations are defined on the open sets in \mathfrak{U} with constant transition functions g_{ij} ; this principal bundle is thus flat by definition. Furthermore, the vertex potential f and the compatibility constraints (1) ensure that the following global section $s : X \rightarrow \mathcal{P}_\rho$ is well-defined on this bundle:

$$s(x) = \phi_i(x, f_i), \quad x \in U_i$$

where $\phi_i : U_i \times G \rightarrow \mathcal{P}_\rho$ is the local trivialization of \mathcal{P}_ρ over U_i . The triviality of this principal bundle then follows from the existence of such a global section; see e.g. [Ste51, §8.3]. The other direction of the proposition follows immediately from this triviality criterion for principal bundles. \square

Proof of Formula (37).

$$\begin{aligned} \langle \omega, \eta \rangle &= \frac{1}{2} \sum_{(i,j) \in E} \left[w_{ij} \langle p_i(\omega_{ij}^{(i)}), p_i(\eta_{ij}^{(i)}) \rangle_F + w_{ji} \langle p_j(\omega_{ji}^{(j)}), p_j(\eta_{ji}^{(j)}) \rangle_F \right] \\ &= \frac{1}{2} \sum_{(i,j) \in E} \left[w_{ij} \langle p_i(\omega_{ij}^{(i)}), p_i(\eta_{ij}^{(i)}) \rangle_F + w_{ij} \langle \rho_{ij} p_j(\omega_{ji}^{(j)}), \rho_{ij} p_j(\eta_{ji}^{(j)}) \rangle_F \right] \\ &= \frac{1}{2} \sum_{(i,j) \in E} \left[w_{ij} \langle p_i(\omega_{ij}^{(i)}), p_i(\eta_{ij}^{(i)}) \rangle_F + w_{ij} \langle p_i(\omega_{ji}^{(i)}), \rho_{ij} p_i(\eta_{ji}^{(i)}) \rangle_F \right] \quad (\text{see compatibility condition (31)}) \\ &= \frac{1}{2} \sum_{(i,j) \in E} \left[w_{ij} \langle p_i(\omega_{ij}^{(i)}), p_i(\eta_{ij}^{(i)}) \rangle_F + w_{ij} \langle p_i(\omega_{ij}^{(i)}), \rho_{ij} p_i(\eta_{ij}^{(i)}) \rangle_F \right] \quad (\text{skew-symmetry}) \\ &= \sum_{(i,j) \in E} w_{ij} \langle p_i(\omega_{ij}^{(i)}), p_i(\eta_{ij}^{(i)}) \rangle_F. \end{aligned}$$

\square

B Graph Laplacian in Discrete Hodge Theory

Define \mathbb{K} -valued 0- and 1-forms on weighted graph $\Gamma = (V, E, w)$ as

$$\Omega^0(\Gamma) := \{f : V \rightarrow \mathbb{K}\}, \quad \Omega^1(\Gamma) := \{\omega : E \rightarrow \mathbb{K} \mid \omega_{ij} = -\omega_{ji} \ \forall (i, j) \in E\},$$

equipped with natural inner products

$$\begin{aligned} \langle f, g \rangle &:= \sum_i d_i \langle f_i, g_i \rangle_{\mathbb{K}}, \quad \forall f, g \in \Omega^0(\Gamma), \\ \langle \omega, \eta \rangle &:= \sum_{(i,j) \in E} w_{ij} \langle \omega_{ij}, \eta_{ij} \rangle_{\mathbb{K}}, \quad \forall \omega, \eta \in \Omega^1(\Gamma), \end{aligned}$$

where $\langle \cdot, \cdot \rangle_{\mathbb{K}}$ is an inner product on \mathbb{K} , and $d_i = \sum_{j:(i,j) \in E} w_{ij}$ is the weighted degree at vertex $i \in V$. Analogous to the study of differential forms on a smooth manifold, one can define the *differential* $d : \Omega^0(\Gamma) \rightarrow \Omega^1(\Gamma)$ and *codifferential* $\delta : \Omega^1(\Gamma) \rightarrow \Omega^0(\Gamma)$ operators that are formal adjoints of each other:

$$(df)_{ij} = f_i - f_j, \quad \forall f \in \Omega^0(\Gamma), \quad (\delta\omega)_i := \frac{1}{d_i} \sum_{j:(i,j) \in E} w_{ij} \omega_{ij}, \quad \forall \omega \in \Omega^1(\Gamma).$$

These constructions can be encoded into a de Rham cochain complex

$$0 \rightleftarrows \Omega^0(\Gamma) \xrightleftharpoons[\delta]{d} \Omega^1(\Gamma) \rightleftarrows 0,$$

which realizes L_0^{rw} , the *graph random walk Laplacian*, as the Hodge Laplacian of degree zero:

$$\left(\Delta^{(0)}f\right)_i := (\delta df)_i = \frac{1}{d_i} \sum_{j:(i,j) \in E} w_{ij} (f_i - f_j) = \left(L_0^{\text{rw}}f\right)_i, \quad \forall i \in V, \quad \forall f \in \Omega^0(\Gamma).$$

It is well known that L_0^{rw} differs from the *normalized graph Laplacian* L_0 by a similarity transform $L_0 = D^{-1/2} L_0^{\text{rw}} D^{1/2}$, where D is a diagonal matrix with weighted degrees of each vertex on its diagonal.

Software MATLAB code implementing SynCut for the numerical simulations and application in automated geometric morphometrics is publicly available at <https://github.com/trgao10/GOS-SynCut>.

Acknowledgement The authors would like to thank Ingrid Daubechies, Robert Bryant, Douglas Boyer, Paweł Dłotko, Kathryn Hess, Vidit Nanda, Shmuel Weinberger, Steve Smale, Lek-Heng Lim, and Pankaj Agarwal for many useful discussions. TG would also like to thank Ingrid Daubechies for her generous and continuous support.

References

- [AADL13] Reema Al-Aifari, Ingrid Daubechies, and Yaron Lipman. Continuous Procrustes Distance Between Two Surfaces. *Communications on Pure and Applied Mathematics*, 66(6):934–964, 2013. [3.1](#), [4.1](#), [4.2](#), [5](#)
- [AB] Dmitri Victorovich Anosov and Andrei Andreevich Bolibruch. The Riemann-Hilbert Problem, *Aspects of Mathematics* 22 (1994). [1.3](#)
- [AB83] Michael Francis Atiyah and Raoul Bott. The Yang-Mills Equations over Riemann Surfaces. *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, 308(1505):523–615, 1983. [2.1](#)
- [AFW06] Douglas N. Arnold, Richard S. Falk, and Ragnar Winther. Finite Element Exterior Calculus, Homological Techniques, and Applications. *Acta Numerica*, 15:1–155, 2006. [1](#)
- [AHTK99] S. Angenent, S. Haker, A. Tannenbaum, and R. Kikinis. On the Laplace-Beltrami Operator and Brain Surface Flattening. *IEEE Transactions on Medical Imaging*, 18(8):700–711, Aug 1999. [4.1](#)
- [AKPW95] Noga Alon, Richard M. Karp, David Peleg, and Douglas West. A Graph-Theoretic Game and its Application to the k-Server Problem. *SIAM Journal on Computing*, 24(1):78–100, 1995. [2.2.4](#)
- [APL15] Noam Aigerman, Roi Poranne, and Yaron Lipman. Seamless Surface Mappings. *ACM Transactions on Graphics (TOG)*, 34(4):72, 2015. [3.1](#)
- [ASC11] Mathieu Aubry, Ulrich Schlickewei, and Daniel Cremers. The Wave Kernel Signature: A Quantum Mechanical Approach to Shape Analysis. In *Computer Vision Workshops (ICCV Workshops), 2011 IEEE International Conference on*, pages 1626–1633, Nov 2011. [3.1](#)
- [BBK08] Alexander M Bronstein, Michael M Bronstein, and Ron Kimmel. *Numerical Geometry of Non-Rigid Shapes*. Springer Science & Business Media, 2008. [3.1](#)
- [BBL⁺11] Jonathan M Bunn, Doug M Boyer, Yaron Lipman, Elizabeth M St Clair, Jukka Jernvall, and Ingrid Daubechies. Comparing Dirichlet Normal Surface Energy of Tooth Crowns, A New Technique of Molar Shape Quantification for Dietary Inference, With Previous Methods in Isolation and in Combination. *American Journal of Physical Anthropology*, 145(2):247–261, 2011. [4.1](#)
- [BC08] Hans-Jurgen Bandelt and Victor Chepoi. Metric Graph Theory and Geometry: A Survey. *Contemporary Mathematics*, 453:49–86, 2008. [1.1](#)
- [BCS15] Afonso S. Bandeira, Yutong Chen, and Amit Singer. Non-unique Games over Compact Groups and Orientation Estimation in Cryo-EM. *arXiv preprint arXiv:1505.03840*, 2015. [1](#), [1.1](#), [1.3](#), [4.1](#)

- [BCSZ14] Afonso S. Bandeira, Moses Charikar, Amit Singer, and Andy Zhu. Multireference Alignment using Semidefinite Programming. In *Proceedings of the 5th Conference on Innovations in Theoretical Computer Science*, pages 459–470. ACM, 2014. [1](#), [1.3](#), [3.1](#), [3.2](#)
- [BD10] Joseph Blitzstein and Persi Diaconis. A Sequential Importance Sampling Algorithm for Generating Random Graphs with Prescribed Degrees. *Internet Math.*, 6(4):489–522, 2010. [3.3.1](#)
- [BHS11] R. Brown, P.J. Higgins, and R. Sivera. *Nonabelian Algebraic Topology: Filtered Spaces, Crossed Complexes, Cubical Homotopy Groupoids*. EMS series of lectures in mathematics. European Mathematical Society, 2011. [4.1](#)
- [BK10] Michael M. Bronstein and Iasonas Kokkinos. Scale-Invariant Heat Kernel Signatures for Non-rigid Shape Recognition. In *Computer Vision and Pattern Recognition (CVPR), 2010 IEEE Conference on*, pages 1704–1711. IEEE, 2010. [3.1](#)
- [BKS16] Afonso S. Bandeira, Christopher Kennedy, and Amit Singer. Approximating the Little Grothendieck Problem over the Orthogonal and Unitary Groups. *Mathematical Programming*, pages 1–43, 2016. [1](#), [1.1](#), [2.2.1](#), [2.2.3](#), [3.2](#)
- [BLS⁺11] Doug M. Boyer, Yaron Lipman, Elizabeth St. Clair, Jesús Puente, Biren A. Patel, Thomas Funkhouser, Jukka Jernvall, and Ingrid Daubechies. Algorithms to Automatically Quantify the Geometric Similarity of Anatomical Surfaces. *Proceedings of the National Academy of Sciences*, 108(45):18221–18226, 2011. [3.1](#), [4.2](#), [4](#)
- [BN03] Mikhail Belkin and Partha Niyogi. Laplacian Eigenmaps for Dimensionality Reduction and Data Representation. *Neural Comput.*, 15(6):1373–1396, June 2003. [1](#)
- [BN04] Mikhail Belkin and Partha Niyogi. Semi-Supervised Learning on Riemannian Manifolds. *Machine Learning*, 56(1-3):209–239, 2004. [1](#)
- [Bol90] Andrey Andreevich Bolibrukh. The Riemann-Hilbert Problem. *Russian Mathematical Surveys*, 45(2):1, 1990. [1.3](#)
- [Boy08] Doug M Boyer. Relief Index of Second Mandibular Molars is a Correlate of Diet Among Prosimian Primates and Other Euarchontan Mammals. *Journal of Human Evolution*, 55(6):1118–1137, 2008. [4.1](#)
- [BPG⁺15] Doug M. Boyer, Jesús Puente, Justin T. Gladman, Chris Glynn, Sayan Mukherjee, Gabriel S. Yapuncich, and Ingrid Daubechies. A New Fully Automated Approach for Aligning and Comparing Shapes. *The Anatomical Record*, 298(1):249–276, 2015. [3.1](#), [4.1](#)
- [Bro67] Ronald Brown. Groupoids and van Kampen’s Theorem. In *Proc. London Math. Soc.*(3), volume 17, pages 385–401. Citeseer, 1967. [4.1](#)
- [Bry07] Jean-Luc Brylinski. *Loop Spaces, Characteristic Classes and Geometric Quantization*, volume 107. Springer Science & Business Media, 2007. [2.1](#)
- [BSAB14] Nicolas Boumal, Amit Singer, Pierre-Antoine Absil, and Vincent D. Blondel. Cramér–Rao Bounds for Synchronization of Rotations. *Information and Inference*, 3(1):1–39, 2014. [1.1](#), [1.1](#), [1.3](#)
- [BSS13] Afonso S. Bandeira, Amit Singer, and Daniel A. Spielman. A Cheeger Inequality for the Graph Connection Laplacian. *SIAM Journal on Matrix Analysis and Applications*, 34(4):1611–1630, 2013. [1](#), [1.1](#), [1.1](#), [1.1](#), [1.1](#), [1.2.2](#), [1.3](#), [2.2](#), [2.2.1](#), [2.2.3](#), [2.2.4](#), [2.2.4](#), [2.2.4](#), [2.2.4](#), [3.2](#), [3.3.2](#), [3](#)
- [BSSZ08] Alexander I. Bobenko, John M. Sullivan, Peter Schröder, and Günter Ziegler. *Discrete Differential Geometry*. Springer, 2008. [1](#)
- [BT82] Raoul Bott and Loring W. Tu. *Differential Forms in Algebraic Topology*, volume 82 of *Graduate Texts in Mathematics*. Springer-Verlag New York, 1982. [1.1](#), [1.3](#), [2.2.3](#)

- [BtD03] Theodor Bröcker and Tammo tom Dieck. *Representations of Compact Lie Groups*. Graduate Texts in Mathematics. Springer Berlin Heidelberg, 2003. [2.2.3](#)
- [CdGDS13] Keenan Crane, Fernando de Goes, Mathieu Desbrun, and Peter Schröder. Digital Geometry Processing with Discrete Exterior Calculus. In *ACM SIGGRAPH 2013 courses*, SIGGRAPH '13, New York, NY, USA, 2013. ACM. [1](#), [2.2.2](#)
- [CGH14] Yuxin Chen, Leonidas Guibas, and Qixing Huang. Near-Optimal Joint Object Matching via Convex Relaxation. In Tony Jebara and Eric P. Xing, editors, *Proceedings of the 31st International Conference on Machine Learning (ICML-14)*, pages 100–108. JMLR Workshop and Conference Proceedings, 2014. [3.1](#)
- [Che66] Nikolai Nikolaevitch Chentsov. A Systematic Theory of Exponential Families of Probability Distributions. *Theor. Probability Appl.*, 11:425–425, 1966. [1](#)
- [Che70] Jeff Cheeger. A Lower Bound for the Smallest Eigenvalue of the Laplacian. In *Problems in analysis (Papers dedicated to Salomon Bochner, 1969)*, pages 195–199. Princeton Univ. Press, Princeton, N. J., 1970. [2.2.4](#)
- [Chu97] Fan R.K. Chung. *Spectral Graph Theory*. Number 92 in CBMS Regional Conference Series in Mathematics. American Mathematical Society, 1997. [1](#)
- [Chu07] Fan R.K. Chung. Four Proofs for the Cheeger Inequality and Graph Partition Algorithms. In *Proceedings of ICCM*, volume 2, page 378, 2007. [2.2.4](#)
- [CKS15] Kunal N. Chaudhury, Yuehaw Khoo, and Amit Singer. Global Registration of Multiple Point Clouds using Semidefinite Programming. *SIAM Journal on Optimization*, 25(1):468–501, 2015. [2.2.1](#), [2.2.3](#), [3.2](#), [3.2](#), [3](#)
- [CL06] Ronald R. Coifman and Stéphane Lafon. Diffusion Maps. *Applied and Computational Harmonic Analysis*, 21(1):5–30, 2006. Special Issue: Diffusion Maps and Wavelets. [1](#), [1.3](#), [4](#)
- [CLL⁺05a] Ronald R. Coifman, Stéphane Lafon, Ann B. Lee, Mauro Maggioni, Boaz Nadler, Frederick Warner, and Steven W. Zucker. Geometric Diffusions as a Tool for Harmonic Analysis and Structure Definition of Data: Diffusion Maps. *Proceedings of the National Academy of Sciences of the United States of America*, 102(21):7426–7431, 2005. [1](#), [1.3](#)
- [CLL⁺05b] Ronald R. Coifman, Stéphane Lafon, Ann B. Lee, Mauro Maggioni, Boaz Nadler, Frederick Warner, and Steven W. Zucker. Geometric Diffusions as a Tool for Harmonic Analysis and Structure Definition of Data: Multiscale Methods. *Proceedings of the National Academy of Sciences of the United States of America*, 102(21):7432–7437, 2005. [1](#), [1.3](#)
- [CLS12] Mihai Cucuringu, Yaron Lipman, and Amit Singer. Sensor Network Localization by Eigenvector Synchronization Over the Euclidean Group. *ACM Transactions on Sensor Networks (TOSN)*, 8(3):19, 2012. [1](#)
- [CLV03] Fan R.K. Chung, Linyuan Lu, and Van Vu. Spectra of random graphs with given expected degrees. *Proceedings of the National Academy of Sciences*, 100(11):6313–6318, 2003. [3.3.1](#)
- [CMOP11] Ozan Candogan, Ishai Menache, Asuman Ozdaglar, and Pablo A. Parrilo. Flows and Decompositions of Games: Harmonic and Potential Games. *Math. Oper. Res.*, 36(3):474–503, August 2011. [1.3](#)
- [COL06] Amin Coja-Oghlan and André Lanka. *The Spectral Gap of Random Graphs with Given Expected Degrees*, pages 15–26. Springer Berlin Heidelberg, Berlin, Heidelberg, 2006. [3.3.1](#)
- [Con85] Alain Connes. Non-commutative Differential Geometry. *Publications Mathématiques de l'Institut des Hautes Études Scientifiques*, 62(1):41–144, 1985. [1](#)
- [Con00] Alain Connes. Noncommutative Geometry. In *Visions in Mathematics*, pages 481–559. Springer, 2000. [1](#)

- [Cor88] Kevin Corlette. Flat G-Bundles with Canonical Metrics. *J. Diff. Geom.*, 28:361–382, 1988. [1.3](#)
- [Cra46] Herald Cramér. *Mathematical methods of statistics*. Princeton University Press, Princeton, 1946. [1](#)
- [Cuc16] Mihai Cucuringu. Sync-Rank: Robust Ranking, Constrained Ranking and Rank Aggregation via Eigenvector and SDP Synchronization. *IEEE Transactions on Network Science and Engineering*, 3(1):58–79, 2016. [1.3](#)
- [CZCG05] Gunnar Carlsson, Afra Zomorodian, Anne Collins, and Leonidas J Guibas. Persistence Barcodes for Shapes. *International Journal of Shape Modeling*, 11(02):149–187, 2005. [1](#)
- [Del70] Pierre Deligne. Équations Différentielles à Points Singuliers Réguliers. *Lecture Notes in Math.*, 163, 1970. [1.3](#)
- [DG03] David L. Donoho and Carrie Grimes. Hessian eigenmaps: new locally linear embedding techniques for high-dimensional data. *Proceedings of the National Academy of Sciences*, 100:5591–5596, 2003. [1](#)
- [DHLM05] Mathieu Desbrun, Anil N Hirani, Melvin Leok, and Jerrold E Marsden. Discrete Exterior Calculus. *arXiv preprint math/0508341*, 2005. [2.2.2](#)
- [DM98] Ian L Dryden and Kanti V Mardia. *Statistical Shape Analysis*, volume 4. John Wiley & Sons New York, 1998. [4.1](#)
- [Dup76] Johan L. Dupont. Simplicial de Rham Cohomology and Characteristic Classes of Flat Bundles. *Topology*, 15(3):233–245, 1976. [1.3](#)
- [EH10] Herbert Edelsbrunner and John Harer. *Computational Topology: An Introduction*. American Mathematical Soc., 2010. [1](#)
- [Esn88] Hélène Esnault. Characteristic Classes of Flat Bundles. *Topology*, 27(3):323–352, 1988. [1.1](#), [1.3](#)
- [EW16] Noureddine El Karoui and Hau-Tieng Wu. Graph Connection Laplacian Methods Can Be Made Robust to Noise. *Ann. Statist.*, 44(1):346–372, 02 2016. [2.2.4](#)
- [EWFJ07] Alistair R Evans, Gregory P Wilson, Mikael Fortelius, and Jukka Jernvall. High-Level Similarity of Dentitions in Carnivorans and Rodents. *Nature*, 445(7123):78–81, 2007. [4.1](#)
- [FAS16] Michaël Fanuel, Carlos M. Alaíz, and Johan A. K. Suykens. Magnetic Eigenmaps for Community Detection in Directed Networks. *ArXiv e-prints*, June 2016. [1.3](#)
- [Fis22] Ronald Aylmer Fisher. On the mathematical foundations of theoretical statistics. *Phil. Trans. Roy. Soc. London, A*, 222:309–368, 1922. [1](#)
- [FL90] Yves Félix and René Lavendhomme. On de Rham’s Theorem in Synthetic Differential Geometry. *Journal of Pure and Applied Algebra*, 69(1):21–31, 1990. [1](#)
- [FS15] Michaël Fanuel and Johan A. K. Suykens. Deformed Laplacians and spectral ranking in directed networks. *ArXiv e-prints*, November 2015. [1.3](#)
- [Gao15] Tingran Gao. *Hypoelliptic Diffusion Maps and Their Applications in Automated Geometric Morphometrics*. PhD thesis, Duke University, 2015. [1](#), [3.1](#), [4.1](#)
- [Gao16] Tingran Gao. The Diffusion Geometry of Fibre Bundles. *submitted*, 2016. arXiv:1602.02330. [1.3](#), [4.1](#)
- [GBAD⁺16] Paula N Gonzalez, Jimena Barbeito-Andrés, Lucas A D’Addona, Valeria Bernal, and S Ivan Perez. Technical Note: Performance of Semi- and Fully Automated Approaches for Registration of 3D Surface Coordinates in Geometric Morphometric Studies. *American Journal of Physical Anthropology*, 2016. [4.1](#)
- [GD04] John C. Gower and Garnt B. Dijkstra. *Procrustes Problems*, volume 3 of *Oxford Statistical Science Series*. Oxford University Press Oxford, 2004. [4.1](#)

- [Gol82] William M. Goldman. Characteristic Classes and Representations of Discrete Subgroups of Lie Groups. *Bulletin of the American Mathematical Society*, 6(1), 1982. [1.3](#)
- [Gol06] William M. Goldman. Mapping Class Group Dynamics on Surface Group Representations. In *Problems On Mapping Class Groups and Related Topics*, volume 74 of *Proc. Sympos. Pure Math.*, pages 189–214. Amer. Math. Soc., 2006. [1.3](#)
- [Gow75] John C. Gower. Generalized Procrustes Analysis. *Psychometrika*, 40(1):33–51, 1975. [4.1](#)
- [GRR15] Alberto García-Raboso and Steven Rayan. Introduction to Nonabelian Hodge Theory. In *Calabi-Yau Varieties: Arithmetic, Geometry and Physics*, pages 131–171. Springer, 2015. [2.2](#)
- [GYDB] Tingran Gao, Gabriel S. Yapuncich, Ingrid Daubechies, and Doug M. Boyer. Automated Techniques for Comparing Shapes in a Biological Comparative Dataset with High Interspecific Variation: Trade-offs, Limitations, and Progress. Submitted. [3.1](#)
- [Hae91] André Haefliger. Complexes of Groups and Orbihedra. In *Group theory from a geometrical viewpoint. Proceedings of a workshop, held at the International Centre for Theoretical Physics in Trieste, Italy, 26 March to 6 April 1990*, pages 504–540. Singapore: World Scientific, 1991. [4.1](#)
- [Hae92] André Haefliger. Extension of Complexes of Groups. *Ann. Inst. Fourier, Grenoblé*, 42(1-2):275–311, 1992. [4.1](#)
- [HG13] Qixing Huang and Leonidas Guibas. Consistent Shape Maps via Semidefinite Programming. *Computer Graphics Forum, Proc. Eurographics Symposium on Geometry Processing (SGP)*, 32(5):177–186, 2013. [1.1](#), [1.3](#), [3.1](#)
- [Hit87] Nigel J. Hitchin. The Self-Duality Equations on a Riemann Surface. *Proc. London Math. Soc.*, 55(3):59–126, 1987. [1.3](#)
- [Hit90] Nigel J. Hitchin. Flat Connections and Geometric Quantization. *Comm. Math. Phys.*, 131(2):347–380, 1990. [2.1](#)
- [HKP12] Christopher Hoffman, Matthew Kahle, and Elliot Paquette. Spectral Gaps of Random Graphs and Applications to Random Topology. *arXiv preprint arXiv:1201.0425*, 2012. [3.3.1](#)
- [HTDL13] Richard Hartley, Jochen Trumpf, Yuchao Dai, and Hongdong Li. Rotation Averaging. *International Journal of Computer Vision*, 103(3):267–305, 2013. [1.3](#)
- [HWG14] Qixing Huang, Fan Wang, and Leonidas Guibas. Functional Map Networks for Analyzing and Exploring Large Shape Collections. *ACM Transactions on Graphics (TOG)*, 33(4):36, 2014. [1.3](#)
- [HZG⁺12] Qixing Huang, Guoxin Zhang, Lin Gao, Shimin Hu, Adrian Bustcher, and Leonidas Guibas. An Optimization Approach for Extracting and Encoding Consistent Maps in a Shape Collection. *ACM Transactions on Graphics*, 31:125:1–125:11, 2012. [3.1](#)
- [JG13] James L Johnson and Tom Goldring. Discrete Hodge Theory on Graphs: A Tutorial. *Computing in Science & Engineering*, 15(5):42–55, 2013. [2.2.2](#), [2.2.3](#)
- [JLYY11] Xiaoye Jiang, Lek-Heng Lim, Yuan Yao, and Yinyu Ye. Statistical Ranking and Combinatorial Hodge Theory. *Mathematical Programming*, 127(1):203–244, 2011. [1](#), [1.3](#), [2.2.2](#), [2.2.3](#)
- [KABL14] Shahar Z Kovalsky, Noam Aigerman, Ronen Basri, and Yaron Lipman. Controlling Singular Values with Semidefinite Programming. *ACM Trans. Graph.*, 33(4):68–1, 2014. [3.1](#)
- [Kas79] Masaki Kashiwara. Faisceaux Constructibles et Systèmes Holonomes d’équations aux Dérivées Partielles Linéaires à Points Singuliers Réguliers. *Séminaire Équations aux dérivées partielles (Polytechnique)*, pages 1–6, 1979. [1.3](#)
- [Kas84] Masaki Kashiwara. The Riemann-Hilbert Problem for Holonomic Systems. *Publications of the Research Institute for Mathematical Sciences*, 20(2):319–365, 1984. [1.3](#)

- [KBLB12] Iasonas Kokkinos, Michael M. Bronstein, Roei Litman, and Alexander M. Bronstein. Intrinsic Shape Context Descriptors for Deformable Shapes. In *Computer Vision and Pattern Recognition (CVPR), 2012 IEEE Conference on*, pages 159–166. IEEE, 2012. [3.1](#)
- [Ken11] Richard Kenyon. Spanning Forests and the Vector Bundle Laplacian. *Ann. Probab.*, 39(5):1983–2017, 2011. [1.1](#), [1.3](#)
- [KH15] Patrice Koehl and Joel Hass. Landmark-Free Geometric Methods in Biological Shape Analysis. *Journal of The Royal Society Interface*, 12(113):20150795, 2015. [3.1](#), [4.2](#)
- [KKBL15] Itay Kezurer, Shahar Z. Kovalsky, Ronen Basri, and Yaron Lipman. Tight Relaxation of Quadratic Matching. *Computer Graphics Forum*, 34(5):115–128, 2015. [1](#), [3.1](#)
- [KLP⁺16] Rasmus Kyng, Yin Tat Lee, Richard Peng, Sushant Sachdeva, and Daniel A. Spielman. Sparsified Cholesky and Multigrid Solvers for Connection Laplacians. In *Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing*, STOC 2016, pages 842–850, New York, NY, USA, 2016. ACM. [2.2.4](#)
- [Koc82] Anders Kock. Differential Forms with Values in Groups. *Bulletin of the Australian Mathematical Society*, 25(03):357–386, 1982. [1](#)
- [Koc06] Anders Kock. *Synthetic Differential Geometry*, volume 333. Cambridge University Press, 2006. [1](#)
- [KT67] Franz W. Kamber and Philippe Tondeur. Flat Bundles and Characteristic Classes of Group-Representations. *American Journal of Mathematics*, 89(4):857–886, 1967. [1.3](#)
- [Lab13] François Labourie. *Lectures on Representations of Surface Groups*. Zurich Lectures in Advanced Mathematics. European Mathematical Society, 2013. [1.3](#)
- [Laf04] Stéphane S Lafon. *Diffusion Maps and Geometric Harmonics*. PhD thesis, Yale University, 2004.
- [LD11] Yaron Lipman and Ingrid Daubechies. Conformal Wasserstein Distances: Comparing Surfaces in Polynomial Time. *Advances in Mathematics*, 227(3):1047–1077, 2011. [3.1](#), [4.1](#), [4.2](#)
- [Lee03] John Marshall Lee. *Introduction to Smooth Manifolds*, volume 218 of *Graduate Texts in Mathematics*. Springer, New York, Berlin, Heidelberg, 2003. [2.2.1](#)
- [Lim15] Lek-Heng Lim. Hodge Laplacians on Graphs. *arXiv preprint arXiv:1507.05379*, 2015. [1](#), [1.3](#), [2.2.2](#), [2.2.3](#)
- [LPD13] Yaron Lipman, Jesús Puente, and Ingrid Daubechies. Conformal Wasserstein Distance: II. Computational Aspects and Extensions. *Math. Comput.*, 82(281), 2013. [3.1](#), [4.1](#), [4.2](#)
- [Lue76] Huei Shyong Lue. Characteristic Classes for the Deformation of Flat Connections. *Transactions of the American Mathematical Society*, 217:379–393, 1976. [1.3](#)
- [LZ14] Rongjie Lai and Hongkai Zhao. Multi-Scale Non-Rigid Point Cloud Registration Using Robust Sliced-Wasserstein Distance via Laplace-Beltrami Eigenmap. *arXiv preprint arXiv:1406.3758*, 2014. [3.1](#), [4.1](#)
- [Mad99] John Madore. *An Introduction to Noncommutative Differential Geometry and Its Physical Applications*, volume 257. Cambridge University Press, 1999. [1](#)
- [Maj13] Shahn Majid. Noncommutative Riemannian Geometry on Graphs. *Journal of Geometry and Physics*, 69:74–93, 2013. [1](#)
- [Mal36] Prasanta Chandra Malhanobis. On the Generalized Distance in Statistics. *Proc. Nat. Inst. Sc. Ind.*, 2:49–55, 1936. [1](#)
- [MDK⁺16] Haggai Maron, Nadav Dym, Itay Kezurer, Shahar Kovalsky, and Yaron Lipman. Point Registration via Efficient Convex Relaxation. *ACM Transactions on Graphics (TOG)*, 35(4):73, 2016. [3.1](#)

- [Meb80] Zoghman Mebkhout. Sur le Probleme de Hilbert-Riemann. In *Complex Analysis, Microlocal Calculus and Relativistic Quantum Theory*, pages 90–110. Springer, 1980. [1.3](#)
- [Meb84] Zoghman Mebkhout. Une Autre Équivalence de catégories. *Compositio Mathematica*, 51(1):63–88, 1984. [1.3](#)
- [MH08] Laurens van der Maaten and Geoffrey Hinton. Visualizing Data using t-SNE. *Journal of Machine Learning Research*, 9(Nov):2579–2605, 2008. [6](#)
- [Mic08] Peter W. Michor. *Topics in Differential Geometry*, volume 93. American Mathematical Soc., 2008. [2.1](#)
- [Mil58] John Milnor. On the Existence of a Connection with Curvature Zero. *Commentarii Mathematici Helvetici*, 32(1):215–223, 1958. [1.3](#)
- [Mor01] Shigeyuki Morita. *Geometry of Characteristic Classes*, volume 199. American Mathematical Soc., 2001. [2](#)
- [MP07] Daniel Martinec and Tomas Pajdla. Robust Rotation and Translation Estimation in Multiview Reconstruction. In *2007 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–8. IEEE, 2007. [1.3](#)
- [MS74] John Milnor and James D. Stasheff. *Characteristic Classes*, volume 76 of *Annals of Mathematics Studies*. Princeton University Press, 1974. [2.2.1](#), [2.2.1](#)
- [MS16] Sayan Mukherjee and John Steenbergen. Random Walks on Simplicial Complexes and Harmonics. *Random Structures & Algorithms*, 2016. [1.3](#), [2.2.4](#)
- [NBCW⁺11a] Andy Nguyen, Mirela Ben-Chen, Katarzyna Welnicka, Yinyu Ye, and Leonidas Guibas. An Optimization Approach to Improving Collections of Shape Maps. *Computer Graphics Forum*, 30(5):1481–1491, 2011. [1.1](#)
- [NBCW⁺11b] Andy Nguyen, Mirela Ben-Chen, Katarzyna Welnicka, Yinyu Ye, and Leonidas Guibas. An Optimization Approach to Improving Collections of Shape Maps. *Computer Graphics Forum*, 30(5):1481–1491, 2011. [1.3](#), [3.1](#)
- [NRV13] Assaf Naor, Oded Regev, and Thomas Vidick. Efficient Rounding for the Noncommutative Grothendieck Inequality. In *Proceedings of the forty-fifth annual ACM symposium on Theory of computing*, pages 71–80. ACM, 2013. [3.2](#)
- [PP93] Ulrich Pinkall and Konrad Polthier. Computing Discrete Minimal Surfaces and Their Conjugates. *Experimental mathematics*, 2(1):15–36, 1993. [4.1](#)
- [PR16] Ori Parzanchevski and Ron Rosenthal. Simplicial Complexes: Spectrum, Homology and Random Walks. *Random Structures & Algorithms*, 2016. [1](#), [1.3](#), [2.2.4](#)
- [PRT15] Ori Parzanchevski, Ron Rosenthal, and Ran J. Tessler. Isoperimetric Inequalities in Simplicial Complexes. *Combinatorica*, pages 1–33, 2015. [1](#), [1.3](#), [2.2.4](#)
- [Pue13] Jesús Puente. *Distances and Algorithms to Compare Sets of Shapes for Automated Biological Morphometrics*. PhD thesis, Princeton University, 2013. [4.1](#)
- [PWM⁺16] James D. Pampush, Julia M. Winchester, Paul E. Morse, Alexander Q. Vining, Douglas M. Boyer, and Richard F. Kay. Introducing molaR: a New R Package for Quantitative Topographic Analysis of Teeth (and Other Topographic Surfaces). *Journal of Mammalian Evolution*, pages 1–16, 2016. [4.1](#)
- [Rao45] Calyampudi Radhakrishna Rao. Information and accuracy obtainable in the estimation of statistical parameters. *Bulletin Calcutta Math. Soc.*, 37:81–91, 1945. [1](#)
- [RBBK10] Dan Raviv, Michael M. Bronstein, Alexander M. Bronstein, and Ron Kimmel. Volumetric Heat Kernel Signatures. In *Proceedings of the ACM workshop on 3D object retrieval*, pages 39–44. ACM, 2010. [3.1](#)

- [RCB97] Anand Rangarajan, Haili Chui, and Fred L Bookstein. The Softassign Procrustes Matching Algorithm. In *Information Processing in Medical Imaging*, pages 29–42. Springer, 1997. [4.1](#)
- [RS00] Sam T. Roweis and Lawrence K. Saul. Nonlinear Dimensionality Reduction by Locally Linear Embedding. *Science*, 290:2323–2326, 2000. [1](#)
- [Sim92] Carlos T. Simpson. Higgs Bundles and Local Systems. *Publications Mathématiques de l’IHÉS*, 75(1):5–95, 1992. [1.3](#)
- [Sim94a] Carlos T. Simpson. Moduli of Representations of the Fundamental Group of a Smooth Projective Variety I. *Publications Mathématiques de l’IHÉS*, 79(1):47–129, 1994. [2.2](#)
- [Sim94b] Carlos T. Simpson. Moduli of Representations of the Fundamental Group of a Smooth Projective Variety. II. *Publications Mathématiques de l’Institut des Hautes Études Scientifiques*, 80(1):5–79, 1994. [2.2](#)
- [Sin11] Amit Singer. Angular Synchronization by Eigenvectors and Semidefinite Programming. *Applied and Computational Harmonic Analysis*, 30(1):20–36, 2011. [1.1](#), [3.2](#)
- [SKM14] John Steenbergen, Caroline Klivans, and Sayan Mukherjee. A Cheeger-Type Inequality on Simplicial Complexes. *Advances in Applied Mathematics*, 56:56–77, 2014. [1](#), [1.3](#), [2.2.4](#)
- [SM00] Jianbo Shi and Jitendra Malik. Normalized Cuts and Image Segmentation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 22(8):888–905, 2000. [3.3.1](#)
- [SNB⁺12] Justin Solomon, Andy Nguyen, Adrian Butscher, Mirela Ben-Chen, and Leonidas Guibas. Soft Maps Between Surfaces. In *Proc. SGP 2012*, 2012. [4.1](#)
- [SOG09] Jian Sun, Maks Ovsjanikov, and Leonidas Guibas. A Concise and Provably Informative Multi-Scale Signature Based on Heat Diffusion. *Computer Graphics Forum*, 28(5):1383–1392, 2009. [3.1](#)
- [SS12] Yoel Shkolnisky and Amit Singer. Viewing Direction Estimation in Cryo-EM using Synchronization. *SIAM Journal on Imaging Sciences*, 5(3):1088–1110, 2012. [1](#), [1.1](#), [1.1](#), [1.1](#), [1.3](#)
- [Ste51] Norman Earl Steenrod. *The Topology of Fibre Bundles*, volume 14. Princeton University Press, 1951. [1.1](#), [1.1](#), [1.1](#), [1.3](#), [2.1](#), [2.1](#), [2.2.1](#), [A](#)
- [SW11] Amit Singer and Hau-Tieng Wu. Orientability and Diffusion Maps. *Applied and Computational Harmonic Analysis*, 31(1):44–58, 2011. [1.1](#), [1.1](#), [1.3](#), [4.1](#)
- [SW12] Amit Singer and Hau-Tieng Wu. Vector Diffusion Maps and the Connection Laplacian. *Communications on Pure and Applied Mathematics*, 65(8):1067–1144, 2012. [1](#), [1.3](#), [2.2](#), [2.2.4](#)
- [SW13] Amit Singer and Hau-tieng Wu. Spectral Convergence of the Connection Laplacian from Random Samples. *arXiv preprint arXiv:1306.1587*, 2013. [1](#), [1.3](#)
- [SZS⁺13] Zhengyu Su, Wei Zeng, Rui Shi, Yalin Wang, Jian Sun, and Xianfeng Gu. Area Preserving Brain Mapping. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2235–2242, 2013. [4.1](#)
- [SZSH11] Amit Singer, Zhizhen Zhao, Yoel Shkolnisky, and Ronny Hadani. Viewing Angle Classification of Cryo-Electron Microscopy Images using Eigenvectors. *SIAM Journal on Imaging Sciences*, 4(2):723–759, 2011. [1](#), [1.1](#), [1.1](#), [1.3](#), [2.2.1](#), [3.1](#)
- [Tau11] Clifford Taubes. *Differential Geometry: Bundles, Connections, Metrics and Curvature*, volume 23. Oxford University Press, 2011. [1.3](#), [2.1](#), [2.1](#), [2.2.3](#)
- [TdSL00] Joshua B. Tenenbaum, Vin de Silva, and John Langford. A Global Geometric Framework for Nonlinear Dimensionality Reduction. *Science*, 290:2319–2323, 2000. [1](#)

- [TSR11] T. Tzveneva, Amit Singer, and Szymon Rusinkiewicz. Global Alignment of Multiple 3-D Scans using Eigenvector Synchronization (bachelor thesis). Technical report, Technical Report, Princeton University, 2011. [1.1](#), [1.1](#), [1.3](#), [3.1](#)
- [Tu83] Loring W. Tu. *Hodge Theory and the Local Torelli Problem*, volume 279. American Mathematical Soc., 1983. [1](#)
- [Vas83] Efsthios Vassiliou. Flat Bundles and Holonomy Homomorphisms. *manuscripta mathematica*, 42(2-3):161–170, 1983. [1.3](#)
- [Vil03] Cédric Villani. *Topics in Optimal Transportation*. Graduate studies in mathematics. American Mathematical Society, 2003. [4.1](#)
- [Vil08] Cédric Villani. *Optimal Transport: Old and New*. Grundlehren der mathematischen Wissenschaften. Springer, 2009 edition, November 2008. [4.1](#)
- [Von07] Ulrike Von Luxburg. A Tutorial on Spectral Clustering. *Statistics and computing*, 17(4):395–416, 2007. [3.3.1](#)
- [Wei95] Alan Weinstein. The Symplectic Structure on Moduli Space. In *The Floer Memorial Volume*, pages 627–635. Springer, 1995. [2.1](#)
- [Wel07] Raymond O. Wells. *Differential Analysis on Complex Manifolds*, volume 65 of *Graduate Texts in Mathematics*. Springer, 2007. [2.1](#)
- [WHG13] Fan Wang, Qixing Huang, and Leonidas Guibas. Image Co-Segmentation via Consistent Functional Maps. In *Computer Vision (ICCV), 2013 IEEE International Conference on*, pages 849–856. IEEE, 2013. [1.3](#)
- [WS13] Lanhui Wang and Amit Singer. Exact and Stable Recovery of Rotations for Robust Synchronization. *Information and Inference*, 2013. [1](#), [1.1](#), [1.1](#), [1.3](#), [2.2.1](#)
- [Xia12] Eugene Z Xia. Abelian and Non-Abelian Cohomology. *Geometry, Topology and Dynamics of Character Varieties*, 23:309–349, 2012. [2.2](#)
- [YL16] Ke Ye and Lek-Heng Lim. Cohomology of Cryo-Electron Microscopy. *arXiv:1604.01319*, 2016. [1.1](#), [1.3](#)
- [ZHT03] Lei Zhu, Steven Haker, and Allen Tannenbaum. Area-Preserving Mappings for the Visualization of Medical Structures. In RandyE. Ellis and TerryM. Peters, editors, *Medical Image Computing and Computer-Assisted Intervention - MICCAI 2003*, volume 2879 of *Lecture Notes in Computer Science*, pages 277–284. Springer Berlin Heidelberg, 2003. [4.1](#)
- [ZSG⁺13] Xin Zhao, Zhengyu Su, X.D. Gu, A Kaufman, Jian Sun, Jie Gao, and Feng Luo. Area-Preservation Mapping using Optimal Mass Transport. *Visualization and Computer Graphics, IEEE Transactions on*, 19(12):2838–2847, Dec 2013. [4.1](#)